



MOX-Report No. 12/2022

**Machine Learning based refinement strategies for
polyhedral grids with applications to Virtual Element
and polyhedral Discontinuous Galerkin methods**

Antonietti, P.F.; Dassi, F.; Manzuzzi, E.

MOX, Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

mox-dmat@polimi.it

<http://mox.polimi.it>

Machine Learning based refinement strategies for polyhedral grids with applications to Virtual Element and polyhedral Discontinuous Galerkin methods

P. F. Antonietti*, F. Dassi†, E. Manuzzi*

February 25, 2022

Abstract

We propose two new strategies based on Machine Learning techniques to handle polyhedral grid refinement, to be possibly employed within an adaptive framework. The first one employs the k-means clustering algorithm to partition the points of the polyhedron to be refined. This strategy is a variation of the well known Centroidal Voronoi Tessellation. The second one employs Convolutional Neural Networks to classify the “*shape*” of an element so that “ad-hoc” refinement criteria can be defined. This strategy can be used to enhance existing refinement strategies, including the k-means strategy, at a low online computational cost. We test the proposed algorithms considering two families of finite element methods that support arbitrarily shaped polyhedral elements, namely the Virtual Element Method (VEM) and the Polygonal Discontinuous Galerkin (PolyDG) method. We demonstrate that these strategies do preserve the structure and the quality of the underlying grids, reducing the overall computational cost and mesh complexity.

Keywords: polyhedral grid refinement, Machine Learning, Convolutional Neural Networks, k-means, Polyhedral Discontinuous Galerkin, Virtual Element Method.

Abbreviations: Machine Learning (ML), Centroidal Voronoi Tessellation (CVT), Convolutional Neural Networks (CNNs), Finite Element Methods (FEMs), Virtual Element Method (VEM), Polyhedral Discontinuous Galerkin (PolyDG).

*MOX, Department of Mathematics, Politecnico di Milano, p.zza Leonardo da Vinci 32, I-20133, Milano, Italy (paola.antonietti@polimi.it, enrico.manuzzi@polimi.it).

†Department of Mathematics and Applications, University of Milano-Bicocca, Via Cozzi 53, I-20153, Milano, Italy (franco.dassi@unimib.it)

1 Introduction

Many applications in the fields of Engineering and Applied Sciences, such as fluid-structure interaction problems, flow in fractured porous media, crack and wave propagation problems, are characterized by a strong complexity of the physical domain, possibly involving moving geometries, heterogeneous media, immersed interfaces (such as e.g. fractures) and complex topographies. Whenever classical Finite Element Methods (FEMs) are employed to discretize the underlying differential model, the process of grid generation can be the bottleneck of the whole simulation, as computational meshes can be composed only of tetrahedral, hexahedral, or prismatic elements. To overcome this limitation, in the last years there has been a great interest in developing FEMs that can employ general polygons and polyhedra as grid elements for the numerical discretizations of partial differential equations. We mention the mimetic finite difference method [1, 2, 3, 4], the hybridizable discontinuous Galerkin method [5, 6, 7, 8], the Polyhedral Discontinuous Galerkin (PolyDG) method [9, 10, 11, 12, 13, 14, 15], the Virtual Element Method (VEM) [16, 17, 18, 19, 20, 21] and the Hybrid High-Order method [22, 23, 24, 25, 26]. This calls for the need to develop effective algorithms to handle polygonal and polyhedral grids and to assess their quality (see e.g. [27]). For a comprehensive overview we refer to the monographs and special issues [4, 14, 25, 28, 20, 21].

Among the open problems, there is the issue of efficiently handling polytopic mesh refinement [29, 30, 31], i.e., partitioning mesh elements into smaller elements to produce a finer grid, and agglomeration strategies, i.e., merging mesh elements to obtain coarser grids [32, 33, 10]. Indeed, as for standard triangular and quadrilateral meshes, during either refinement or agglomeration it is important to preserve the quality of the underlying mesh, since this might affect the overall performance of the method in terms of stability and accuracy. Indeed a suitable adapted mesh may allow to achieve the same accuracy with a much smaller number of degrees of freedom when solving the numerical problem, hence saving memory and computational power. However, since in such a general framework mesh elements may have any shape, there are not well established strategies to achieve effective refinement or agglomeration with a fast, robust and simple approach, contrary to classical tetrahedral, hexahedral and prismatic meshes. Moreover, grid agglomeration is a topic quite unexplored, because it is not possible to develop such kind of strategies within the framework of classical FEMs.

In recent years there has been a great development of Machine Learning (ML) algorithms to enhance and accelerate numerical methods for scientific computing. Examples include, but are not limited to, [34, 35, 36, 37, 38, 39, 40, 41]. In this work, we propose different strategies to handle polyhedral grid refinement, that exploit ML techniques to extract information about the “shape” of mesh elements in order to refine them accordingly. Starting from the approach developed in [42] in two dimensions, here we address the three-dimensional case. In particular, we employ the k-means clustering algorithm [43, 44] to partition

the points of the polyhedron to be refined. Learning a clustered representation allows to better handle situations where elements have no particular structure, while preserving the quality of the grid. This strategy is a variation of the well known Centroidal Voronoi Tessellation [45, 46]. We also explore how to enhance existing refinement strategies, including the k-means, using Convolutional Neural Networks (CNNs), powerful function approximators successfully employed in many areas of ML, especially computer vision [47, 48]. We show that CNNs can be employed to identify correctly the “shape” of a polyhedron, without resorting to the explicit calculation of geometric properties which may be expensive in three dimensions. This information can then be exploited to design tailored strategies for different families of polyhedra. According to the classification, the most suitable refinement strategy is then selected among the available ones, therefore enhancing performance. The proposed approach has several advantages:

- it helps preserving structure and quality of the mesh, since it can be easily tailored for different types of elements;
- it can be combined with suitable user-defined refinement criteria, including the k-means strategy;
- it is independent of the differential model at hand and of the numerical method employed for the discretization;
- the overall computational cost is kept low, since CNNs can be trained offline once and for all.

To investigate the capabilities of the proposed approaches, we consider a second-order model problem discretized by either the VEM and the PolyDG method. We measure effectiveness through an analysis of quality metrics and accuracy of the discretization error.

The paper is organized as follows. In Section 2 we present the k-means refinement strategy and other possible refinement criteria for polyhedra. In Section 3 we propose a general framework to enhance existing refinement strategies using CNNs. In Section 4 we validate the proposed refinement strategies on a set of polyhedral meshes and we measure their effectiveness using the quality metrics introduced in [27]. In Section 5 we present some computations obtained with a uniform and a priori adaptive mesh refinement process for a differential model approximated either by VEMs and PolyDG methods. In Section 6 we draw some conclusions.

2 Refinement Strategies

A polyhedral mesh can be generated, e.g., by standard triangulation algorithms or by hexahedral elements in the structured case. Another possibility is to use a Voronoi Tessellation [30]: suitable points, called seeds, are chosen inside the

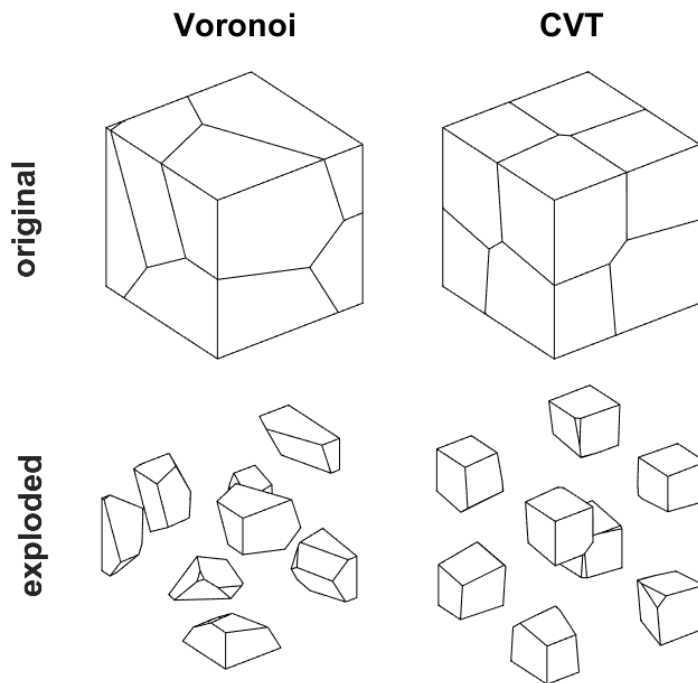


Figure 1: A Voronoi tessellation (top left) and a CVT (top right) of a cube and their exploded versions (bottom). “Small” edges and faces are generated.

polyhedron and each element of the new partition is the set of points which are closer to a specific seed (see Figure 1 left column). A Centroidal Voronoi Tessellation (CVT) [45, 46] is a Voronoi tessellation whose seeds are the centroids (centers of mass) of the corresponding new elements (see Figure 1 right column). This strategy produces elements which are shape-regular and of similar size (these properties will be more formally addressed in Section 4). However, it has a considerable computational cost and it may produce elements with many small edges and faces, which implies higher costs in terms of memory storage and algorithm complexity. Moreover, although small edges and faces do not necessarily deteriorate the accuracy of numerical methods, they are in general not beneficial [49, 50, 51, 52]. As we will see in the following, these algorithms for grids generation can be adapted to perform mesh refinement.

In order to refine a three dimensional polyhedron a possible strategy is to subdivide the polyhedron along a chosen preferential direction [31]. It has a low computational cost and allows to adjust the cutting plane at each iteration. For instance it is possible to choose such plane so that we avoid small edges and consequently yield simpler mesh elements. In particular, we propose the general refinement strategy described in Algorithm 1, where the “size” or “diameter”

of a polyhedron $P \subset \mathbb{R}^3$ is defined, as usual, as

$$\text{diam}(P) = \sup\{\|x - y\|, x, y \in P\},$$

where $\|\cdot\|$ is the standard Euclidean distance. We also recall that given a polyhedral mesh, i.e., a set of non-overlapping polyhedra $\{P_i\}_{i=1}^{N_P}$, $N_P \geq 1$ that covers a domain Ω , the mesh size is defined as

$$h = \max_{i=1, \dots, N_P} \text{diam}(P_i).$$

Possible ways of choosing the cutting plane in step 1 will be discussed in Section 2.1. The “validity check” in step 3, to see if it is possible to slice the element

Algorithm 1: General refinement strategy of a polyhedron P

Input: polyhedral element P to refine, plane tolerance `tol`, maximum number of elements `nmax`, target size \bar{h} .

Output: partition of P into polyhedral sub-elements.

- 1 Choose a plane along which to slice the element.
 - 2 If there are vertices which closer than `tol` to the plane, adjust the plane so that it passes from them.
 - 3 If the plane passes the “validity check”, slice the element into two sub-elements, otherwise use the “emergency strategy”.
 - 4 Until there are less than `nmax` elements, repeat the above steps for each element with a size above \bar{h} .
-

along the chosen plane, may encompass several requirements.

Geometrical. Avoid generating elements with holes and other degenerations, which may occur if the polyhedron is non-convex (see Figure 2).

Numerical. Avoid generating small edges and faces according to a prescribed threshold, which may still occur because in step 2 a plane can be adjusted to pass from 3 points.

If the chosen plane does not pass the “validity check” a possible “emergency strategy” is the following: consider small perturbations of position and orientation of the original plane, until a valid configuration is found.

2.1 Choosing the cutting direction

Several strategies can be designed by making different choices of the cutting direction in step 1 of Algorithm 1. In Figure 3 the proposed strategies are applied on a tetrahedron. Notice that which strategy is the most effective may depend on several factors, such as the initial geometry of the mesh, the quality of the refined elements, the computational cost, the stability of the numerical scheme, or the regularity of the solution.

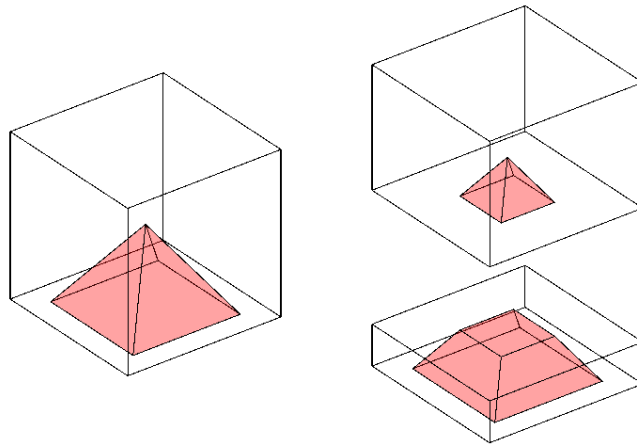


Figure 2: Left: a non-convex polyhedron is obtained from a cube (transparent) by removing a pyramid (red). Right: the polyhedron is sliced along a plane generating two polyhedra, one of which has a hole.

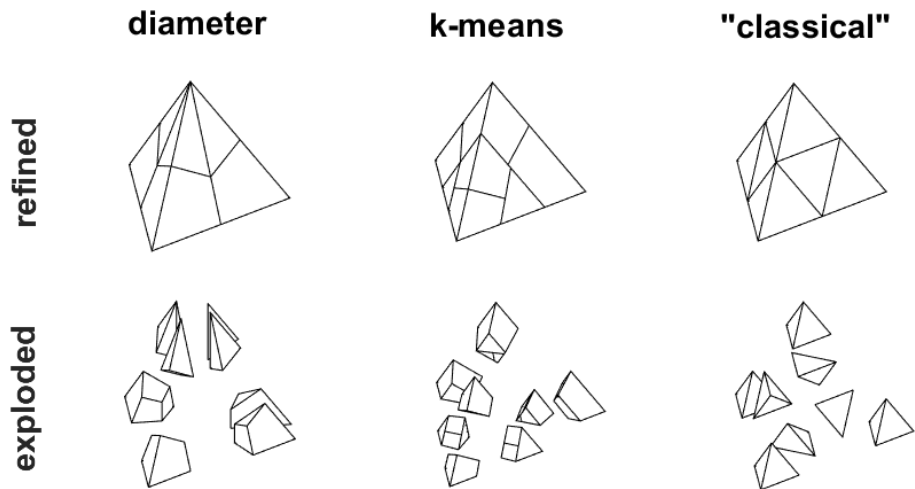


Figure 3: A tetrahedron is refined using diameter, k-means and "classical" strategies (top row) and their exploded versions (bottom row).

Diameter strategy. The diameter of a polyhedron is identified by the two farthest apart vertices of the element, say v_1 and v_2 . The cutting plane has origin $x_0 = (v_1 + v_2)/2$ and is orthogonal to the direction of the diameter, i.e., it has normal $n = (v_2 - v_1)/\|v_2 - v_1\|$. This is a greedy algorithm to halve the size of the element. It has the advantage of easily computing the cutting plane, but it may produce skewed elements because the overall shape of the polyhedron is not taken into account.

K-means strategy. Points belonging to the polyhedron are grouped into two clusters using the k-means algorithm [43, 44] and which are then separated by the cutting plane, as described in Algorithm 2. This is equivalent to a CVT with only two seeds. This strategy produces rounded elements of similar size, but it has a higher computational cost. In principle, different clustering algorithms may be used to learn different representations of the element, such as hierarchical clustering [53, 54] or self-organizing maps [55, 56].

Algorithm 2: K-means cutting plane

Input: polyhedron P to refine, number of grid points n .

Output: cutting plane of P .

- 1 Construct a grid of n of points inside the smallest domain of the form $[x_1, x_2] \times [y_1, y_2] \times [z_1, z_2]$ which contains P .
 - 2 Remove the points of the grid outside P [57], which can be done efficiently when P is convex.
 - 3 Randomly sample two points and label them centroids c_1 and c_2 .
 - 4 Assign each point to the cluster with the closest centroid in L^2 norm.
 - 5 Compute the average of points in each cluster to obtain two new centroid c_1 and c_2 .
 - 6 Repeat steps 4 and 5 until cluster assignments do not change, or the maximum number of iterations is reached.
 - 7 The cutting plane has normal $n = (c_2 - c_1)/\|c_2 - c_1\|$ and origin $x_0 = (c_1 + c_2)/2$.
-

“Classical” strategies. If the shape of the polyhedron is known (up to rotation, scaling, stretching or reflection) it is possible to design tailored refinement strategies using multiple cutting planes. This is the case of tetrahedra, cubes and prisms, which are employed in classical FEMs. These refinement strategies, shown in Figure 4, produce regular elements at low computational cost. Moreover, since the new elements have the same shape of the original one, the algorithm can be applied recursively. Strategies for other polyhedral element can be designed. The main drawback is that these strategies cannot be applied if the shape of the polyhedron is unknown.

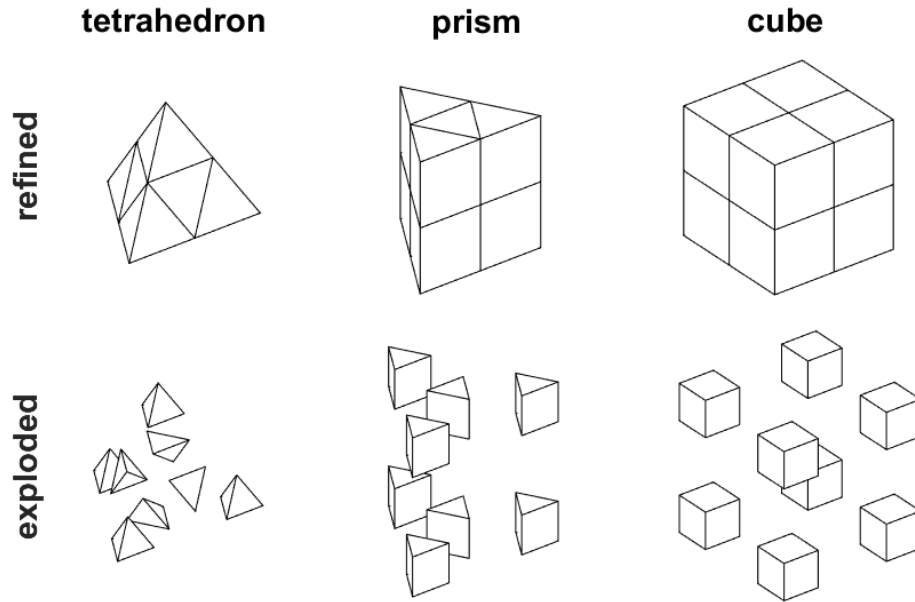


Figure 4: “Classical” refinement strategies for a tetrahedron, a prism and a cube (top row) and their exploded version (bottom row), employed in classical FEMs.

3 Enhancing refinement strategies using CNNs

Consider elements at the top row of Figure 5 from left to right:

1. A tetrahedron where a corner has been cut. This situation may arise when slicing a background mesh, e.g., in order to model a fracture or a soil discontinuity.
2. A prism where one of its faces is divided into two, due to the presence of non-matching grids. This situation can arise each time a neighbouring mesh element is refined, because this may require to partition a face which is shared by two elements.
3. An element of a CVT. These tessellations are employed for meshing irregular domains with few regular elements.
4. A skewed element with no particular structure. These elements can be found in standard Voronoi tessellations, or may appear even after refining a regular element.

Despite the fact that the geometrical representations of elements 1-3 in Figure 5 do not match any classical polyhedron, their overall shapes are similar to a tetrahedron, a prism and a cube, respectively. Therefore refining them with the

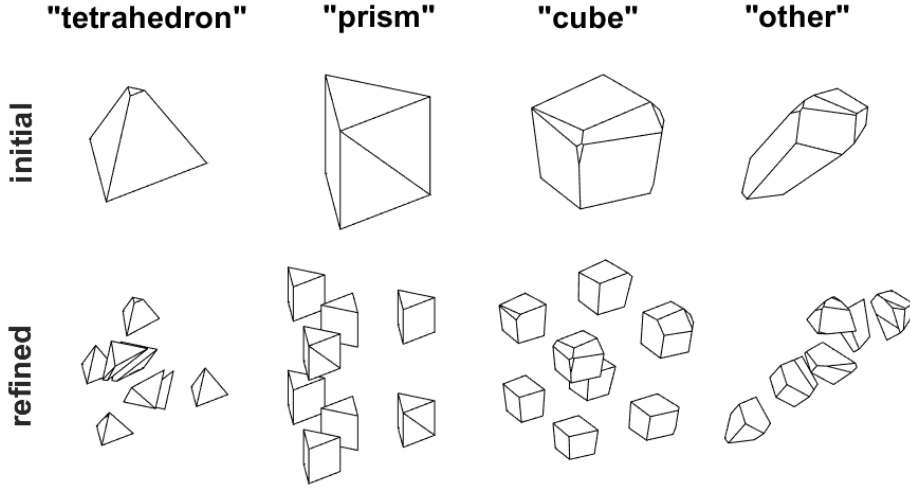


Figure 5: Top row: four polyhedra are classified as “tetrahedron”, “prism”, “cube” and “other”, according to their overall shape and not to their geometrical description. Bottom row: the polyhedra are refined with the corresponding “classical” strategies according to their labels (columns 1-3) and the k-means strategy (column 4).

corresponding shape strategies would produce regular elements at low computational cost. This can be done by computing cutting directions on a “reference shape” which can be constructed using Algorithm 3. The last element does not have a particular shape, and therefore we apply the k-means strategy, which is more robust and flexible even though more expensive. The refined elements are shown at the bottom row of Figure 5. However, in order to apply the correct strategy for each element, we need an algorithm to access whether the shape of a polyhedron is more similar to a tetrahedron, a prism, a cube or none of these. We will use labels “tetrahedron”, “prism”, “cube” and “other”, respectively, and we will refer to such labels as “equivalence classes”. More in general, given a set of polyhedra to refine, $\mathcal{P} = \{P_1, P_2, \dots\}$, and a set of possible refinement strategies, $\mathcal{R} = \{R_1, R_2, \dots, R_n\}$, we want to find a classifier $F : \mathcal{P} \rightarrow \mathcal{R}$ that assigns to every polyhedron the most effective refinement strategy in terms of quality of the refined elements and computational cost. The effectiveness will be measured with suitable metrics later in Section 4.

3.1 Polyhedra classification using CNNs

In principle, any classifier may be used for polyhedra. However, considering geometrical properties of the element alone, such as connectivity and area of the faces, may be too complex to operate a suitable classification: for example, the element labeled as “tetrahedron” in Figure 5 is actually a deformed prism. Approaches like Shape Analysis [60], which relies on applying transformations

Algorithm 3: Computing the reference shape

Input: polyhedron P , shape type S (tetrahedron, cube, prism, ...).

Output: reference shape of P .

- 1 Find n vertices of the polyhedron far apart from each others and from the centroid of the element, using the farthest first traversal algorithm [58], where n is the number of vertices of S .
 - 2 Apply the convex hull algorithm [59] to the n vertices to obtain a triangulated surface.
 - 3 Merge the m couples of neighbouring triangles of the surface whose angle is closest to 180° , where m is the number of quadrilateral faces of S .
-

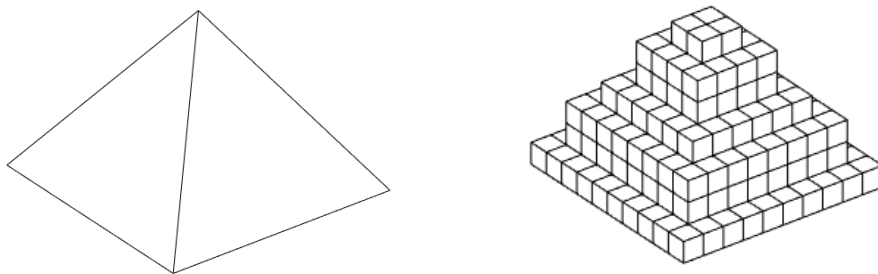


Figure 6: A pyramid (left) and its three-dimensional image (right). Pixels have value 1 inside the polyhedron (represented with cubes) and 0 outside (not represented).

the polyhedron in order to match a reference shape, are too expensive in three-dimensions and seem to lack of the required flexibility: for example there is no reference shape for class “other”. Instead, providing samples of polyhedra together with the desired label can be easily done: for classes “tetrahedron”, “prism” and “cube” we start from the corresponding regular polyhedra and then apply small deformations and rotations, while for class “other” we consider polyhedra from several Voronoi tessellations. This database can be used to “train” our function, i.e., to tune its parameters in order to obtain the desired classifier. This framework, which exploits labeled data, is known as “supervised learning” and CNNs are powerful function approximators in the context of image classification. Indeed, the most intuitive way to asses the shape of an object is by visual inspection and polyhedra can be easily converted into binary images: each pixel assumes value 1 inside the polyhedron and 0 outside [57] (see Figure 6), which can be done efficiently in the convex case. Mathematically, while the geometrical representation of a polyhedron is given by its boundary, the information about the overall shape is given by the distribution of its volume, i.e., the location of its pixels. Notice that a three-dimensional image contains the

same information (up to a scaling and a translation) required by the k-means algorithm, which is a form of “unsupervised” learning because only unlabeled data are provided. The proposed CNN-enhanced refinement strategy is describe in Algorithm 4. In general, the use of CNNs should be seen as tool to enhance

Algorithm 4: CNN-enhanced refinement strategy

Input: polyhedron P to refine.

Output: partition of P into polyhedral sub-elements.

- 1 Construct a binary image of P .
 - 2 Classify the image using a CNN and obtain the label \mathcal{L} .
 - 3 **if** $\mathcal{L} = \text{“tetrahedron”, “prism” or “cube”}$ **then**
 - 4 Apply Algorithm 3 to compute the corresponding reference shape.
 - 5 Compute the cutting directions on the reference shape using the corresponding “classical” strategy.
 - 6 Refine P by applying the cutting directions using Algorithm 1.
 - 7 **end**
 - 8 **if** $\mathcal{L} = \text{“other”}$ **then**
 - 9 Refine P using the k-means strategy, i.e., computing the cutting directions using Algorithm 2 and then using Algorithm 1.
 - 10 **end**
-

existing refinement strategies, not as a refinement strategy on its own in competition with the others. In the case of “classical” strategies, the CNN is extending their domain of applicability to general polyhedra, which would not be possible otherwise.

3.2 Supervised learning for image classification

Consider a three dimensional binary image $\mathbf{B} \in \{0,1\}^{n \times n \times n}$, $n \geq 1$, and the corresponding label vector $y = ([y]_j)_{j=1,\dots,\ell} \in [0,1]^\ell$, where $\ell \geq 2$ is the total number of classes, and $[y]_j$ is the probability of \mathbf{B} to belong to the class j for $j = 1 : \ell$. In a supervised learning framework, we are given a dataset of desired input-output couples $\{(\mathbf{B}_i, y_i)\}_{i=1}^N$, where N is number of labelled data. We consider then an image classifier represented by a function of the form $F : \mathbb{R}^{n \times n \times n} \rightarrow (0,1)^\ell$, in our case a CNN, parameterized by $w \in \mathbb{R}^M$ where $M \geq 1$ is the number of parameters. Our goal is to tune w so that F minimizes the data misfit, i.e.,

$$\min_{w \in \mathbb{R}^M} \sum_{i \in I} l(F(\mathbf{B}_i), y_i),$$

where I is a subset of $\{1, 2, \dots, N\}$ and l is the cross-entropy loss function defined as

$$l(F(\mathbf{B}), y) = \sum_{j=1}^{\ell} -[y]_j \log[F(\mathbf{B})]_j.$$

This optimization phase is also called “learning” or “training phase”. During this phase, a known shortcoming is “overfitting”: the model fits very well the data used in the training phase, but performs poorly on new data. We proceed in a standard way [61], i.e., we split the data into

- *training set*: used to tune the parameters during the training phase;
- *validation set*: used to monitor the model capabilities during the training phase, on data that are not in the training set. The training is halted if the error on the validation set starts to increase;
- *test set*: used to access the actual model performance on new data after the training.

While the training phase can be computationally demanding, because of the large amount of data and parameters to tune, it needs to be performed offline once and for all. Instead, classifying online a new image using a pre-trained model is computationally fast: it requires only to evaluate F on a new input. The predicted label is the one with the highest estimated probability.

3.3 CNN architecture

CNNs are parameterized functions, in our case of the form

$$\text{CNN} : \mathbb{R}^{n \times n \times n} \rightarrow (0, 1)^\ell, \quad n, \ell \geq 2,$$

constructed by composition of simpler functions called “layers of neurons” [47]. We basically use these types of layers.

Convolutional layers. Linear maps performing a convolution of the input, which can be viewed as a filter scanning through the image, extracting local features that depend only on small subregions of the image. This is effective because a key property of images is that close pixels are more strongly correlated than distant ones. The scanning filter mechanism provides the basis for the invariance of the output to translations and distortions of the input image [61]. Multiple convolutions may be performed in a single layer, each of which is called a “feature map”.

Pooling layers. Maps used to perform down-sampling by merging semantically similar features. They perform operations region-wise, such as averaging or computing the maximum. These regions are rectangular and they are computed by fixing a size and then moving along the image of an amount of pixels called “stride”. They improve the invariance of the output with respect to translations of the input.

Activation functions. Maps used to introduce non-linearity, such as the rectified linear unit

$$[\text{RELU}(x)]_i = \max(0, x_i).$$

Dense layers. Generic linear maps of the form $\text{LINEAR} : \mathbb{R}^m \rightarrow \mathbb{R}^\ell$, $m, \ell \geq 1$ defined by parameters to be tuned. They are used to separate image features extracted in the previous layers.

Softmax. We define the function $\text{SOFTMAX} : \mathbb{R}^\ell \rightarrow (0, 1)^\ell$, where $\ell \geq 2$ is the number output classes,

$$[\text{SOFTMAX}(x)]_i = \frac{e^{x_i}}{\sum_{j=1}^{\ell} e^{x_j}}.$$

They are used to assign a probability to each class.

In practise, subsequent application of convolutional, activation and pooling layers may be used to obtain a larger degree of invariance to input transformations such as rotations, distortions, etc...

3.4 CNN training

We have used a database of 22500 binary images of size $16 \times 16 \times 16$ pixels, with labels “tetrahedron”, “prism”, “cube” and “other”. The data were divided into training, validation and test set with ratios 60%-20%-20% respectively. They were generated as follows:

- for classes “tetrahedron”, “prism” and “cube” we generated 6000 images each, starting from the corresponding regular polyhedra, and then randomly applying perturbations, such as rotation, scaling, stretching and reflection;
- for the class “other” we generated 4500 images from elements of several Voronoi tessellations, where seeds have been randomly located.

We generated less images for class “other”, because among all of the randomly generated Voronoi elements, some of them will be actually deformed tetrahedra, prisms and cubes. In these doubtful situations, the CNN will be more likely to classify polyhedra as having a particular shape rather assigning them label “other”. Indeed, the CNN is capable to learn the whole data distribution, including the relative frequency of each class. In this way, we resort to the k-means only when really needed.

Generating data which are representative of the application of interest is the most critical part of the process: the CNN can easily misclassify new samples if they are too different from the training set, even if the accuracy on the test set is extremely high. In three dimensions, the variability of polyhedra is high and generating representative sample is much more complex than in the two-dimensional case [42]. It is important to have a class “other” to which unknown

polyhedra will be assigned, in order to treat them with a robust backup strategy, such as the k-means. This should not discourage from using “classical” strategies, because when employed they are likely to produce regular elements.

The CNN architecture we used is $\text{CNN} : \{0, 1\}^{16 \times 16 \times 16} \rightarrow (0, 1)^4$, where

$$\begin{aligned} \text{CNN} = & \text{CONV}(\bar{f} = 8, m = 8) \rightarrow \text{RELU} \rightarrow \text{POOL}(\bar{f} = 2, s = 2) \rightarrow \\ & \text{CONV}(\bar{f} = 4, m = 8) \rightarrow \text{RELU} \rightarrow \text{POOL}(\bar{f} = 2, s = 2) \rightarrow \\ & \text{CONV}(\bar{f} = 2, m = 8) \rightarrow \text{RELU} \rightarrow \text{POOL}(\bar{f} = 2, s = 2) \rightarrow \\ & \text{LINEAR} \rightarrow \text{SOFTMAX}, \end{aligned}$$

where $\text{CONV}(\bar{f}, s)$ is convolutional layer with filter size \bar{f} (i.e. a window $\bar{f} \times \bar{f} \times \bar{f}$) and m feature maps, $\text{POOL}(\bar{f}, s)$ is an average pooling layer with filter size \bar{f} and stride s , RELU and LINEAR were defined the previous Section 2.2, and for any two compatible functions f_1 and f_2 the arrow notation is defined as $f_1 \rightarrow f_2 = f_2 \circ f_1$. We employed average pooling layers, rather than max pooling layers, because they have the effect to blurry and coarse-grain image features. This has been done in order to focus the CNN on the overall shape of the polyhedron, and not on small scale details and corners. Three sequences of $\text{CONV} - \text{RELU} - \text{POOL}$ layers seem to be sufficient to enforce invariance with respect to rotations of the input polyhedron. The size of the images is $16 \times 16 \times 16$ pixels, i.e., a resolution large enough to apply 3 pooling layers, whose effect is to halve the size of the image across all the dimensions. Such a small resolution allows to quickly generate and classify the image of a polyhedron, which is important in an online setting. Moreover, a smaller CNN is more likely to be robust with respect to samples very different from the training data. It was empirically observed that a larger resolution was not needed for the considered dataset. This is motivated by the fact that the approximation properties of neural networks generally depend on the dimension of the data manifold, not on the dimension of the input space [62]. A larger resolution may be needed to classify polyhedra characterized by smaller scale features, e.g., to correctly distinguish between a dodecahedron and an icosahedron. As shown in Figure 7, the class “other” is reasonably the most misclassified, because it includes polyhedra with very different shapes, some of which are actually deformed tetrahedra, prisms and cubes. Training the CNN using the Adam optimizer [63] takes approximately 15 minutes with MATLAB2020b on a Windows OS 10 Pro 64-bit, CPU NVidia GeForce GTX 1050 Ti and 8244Mb Video RAM.

4 Validation on a set of polyhedral grids

In this section we compare the performance of the proposed algorithms. To evaluate the quality of the refined grids, we employ the following quality metrics introduced in [27]:

- *Uniformity Factor* (UF): ratio between the diameter of an element P and

Confusion Matrix

Output Class	tetrahedron	1162 25.8%	0 0.0%	0 0.0%	33 0.7%	97.2% 2.8%
	prism	0 0.0%	1176 26.1%	0 0.0%	35 0.8%	97.1% 2.9%
	cube	0 0.0%	0 0.0%	1248 27.7%	45 1.0%	96.5% 3.5%
	other	7 0.2%	0 0.0%	0 0.0%	794 17.6%	99.1% 0.9%
		99.4% 0.6%	100% 0.0%	100% 0.0%	87.5% 12.5%	97.3% 2.7%
	tetrahedron	prism	cube	other		
	Target Class					

Figure 7: Confusion matrix for polyhedra classification using a CNN. Class “other” is reasonably the most misclassified, because it includes polyhedra with very different shapes, some of which are actually deformed tetrahedra, prisms and cubes.

the mesh size

$$\text{UF}(P) = \frac{\text{diam}(P)}{h}.$$

This metric takes values in $[0, 1]$. The higher its value is the more mesh elements have comparable sizes.

- *Circle Ratio* (CR): ratio between the radius of the inscribed circle and the radius of the circumscribed circle of a polyhedron P

$$\text{CR}(P) = \frac{\max_{\{B(r) \subset P\}} r}{\min_{\{P \subset B(r)\}} r},$$

where $B(r)$ is a ball of radius r . For the practical purpose of measuring the roundness of an element the radius of the circumscribed circle has been approximated with $\text{diam}(P)/2$. This metric takes values in $[0, 1]$. The higher its value is the more rounded mesh elements are.

Moreover, to measure the computational complexity of the different refinement strategies, for the refined grids we will consider the number of vertices, of edges, of faces, of elements, the total refinement time and the average refinement time per element.

In the following, we apply Algorithm 1 to a polyhedron P using a tolerance $\text{tol} = \text{diam}(P) \cdot 10^{-3}$, a maximum number of elements $\text{nmax} = 8$ (needed to apply the ‘‘classical’’ refinement strategy for cubes), and a target size $\bar{h} = \text{diam}(P)/2$. Concerning the k-means strategy, we apply Algorithm 2 using a grid of $n = 20^3$ points, which is comparable to the size of 16^3 pixels required by the CNN. This value of n may seem high, but it is needed in order compute the cutting plane with a satisfactory accuracy. Indeed, many grid points lie outside the polyhedron and therefore they are not included in the computation of the clusters, which is also the most expensive part of the process. We consider five different coarse grids of domain $(0, 1)^3$: a grid of tetrahedra, a grid of cubes, a grid of prisms, a Voronoi grid with random seeds location, and a CVT. In Figure 8 these grids have been refined uniformly, i.e., each mesh element has been refined, for three times using the diameter strategy (diameter), the k-means strategy (k-means) and the CNN-enhanced strategy (CNN). As we can see, the initial mesh geometry seems to be disrupted by the diameter strategy, approximately preserved by the k-means strategy, and well preserved by the CNN-enhanced strategy.

In Figure 9 we show the computed quality metrics for the uniformly refined grids. As we can see, in general quality metrics are lower for the diameter strategy. The k-means strategy produces elements with higher Circle Ratio, while the CNN-enhanced strategy with higher of Uniformity Factor, although in many cases both of them are comparable.

In Figure 10 we show, for all the considered grids, the number of vertices, edges, faces, elements, the total refinement time and the average refinement time per element. These results have been obtained by using MATLAB2020b on a Windows OS 10 Pro 64-bit, Intel(R) Core(TM) i7-8750H CPU (2.20GHz / 2.21GHz) and 16 GB RAM memory. As we can see, in general the diameter

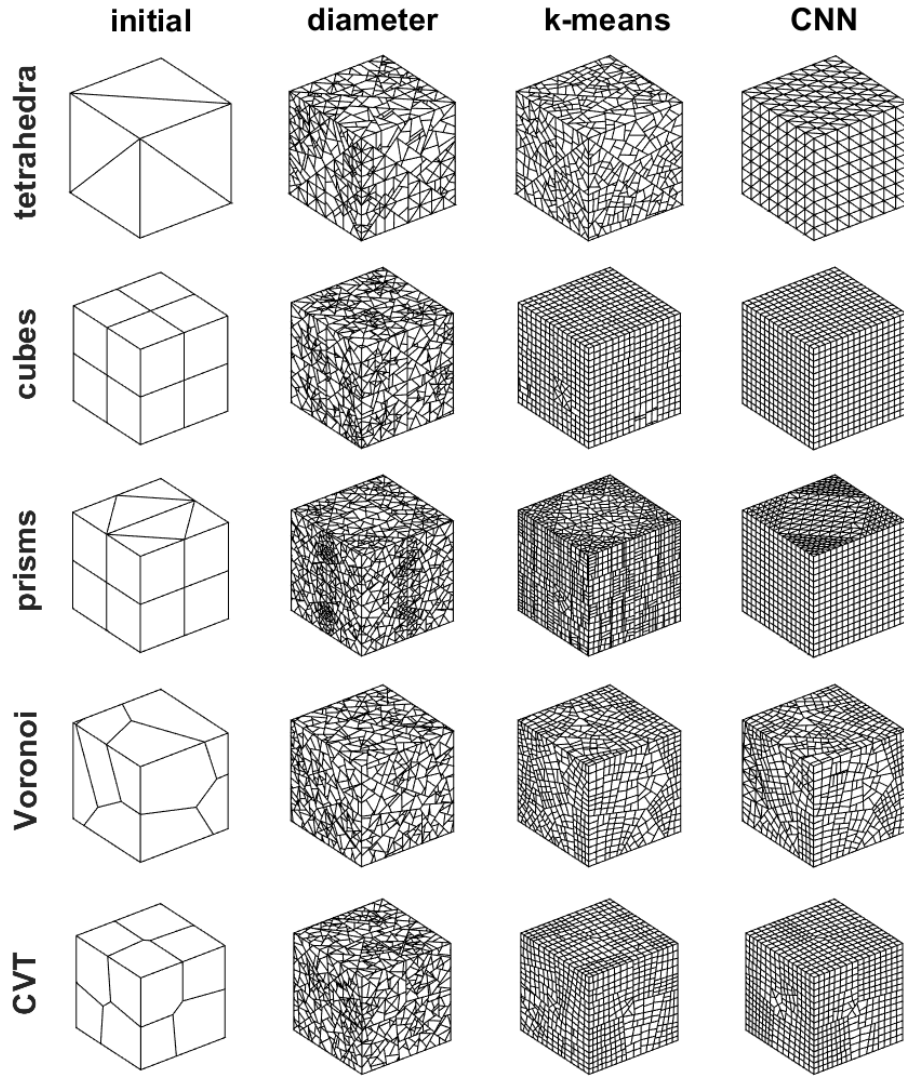


Figure 8: In the first column, coarse grids of domain $(0, 1)^3$: a grid of tetrahedra, a grid of cubes, a grid of prisms, a Voronoi grid with random seeds location, and a CVT. Second to fourth columns: refined grids obtained after three steps of uniform refinement based on employing the k-means strategy (k-means) and the CNN-enhanced strategy (CNN). Each row corresponds to the same initial grid, while each column corresponds to the same refinement strategy.

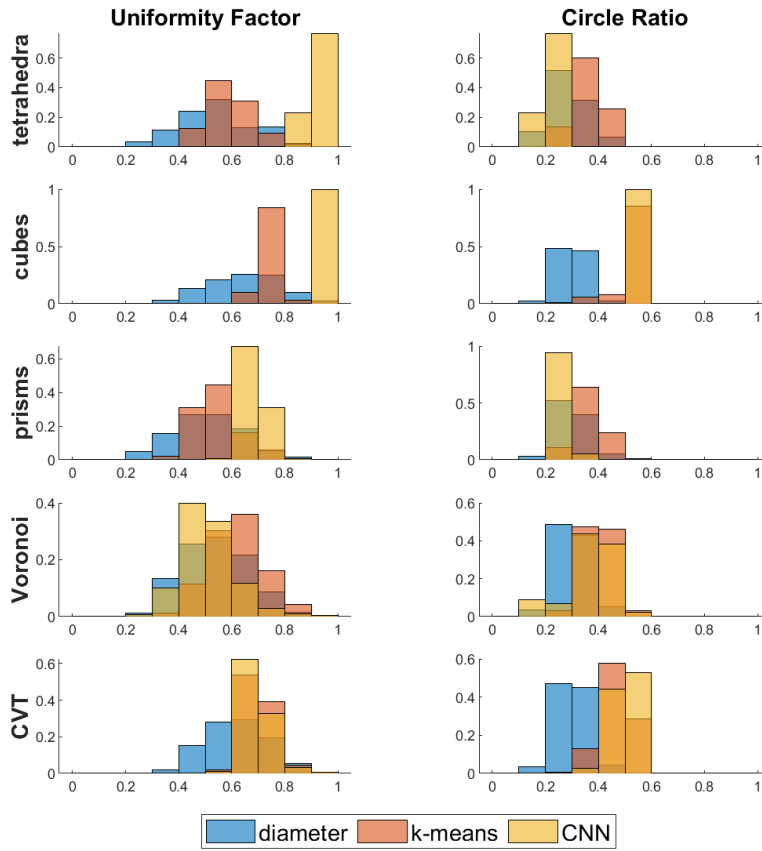


Figure 9: Histograms of the computed quality metrics (Uniformity Factor and Circle Ratio on the x-axis and percentage of grid elements on the y-axis) for the refined grids reported in Figure 8 (second to fourth column), obtained based on employing different refinement strategies (diameter, k-means, CNN).

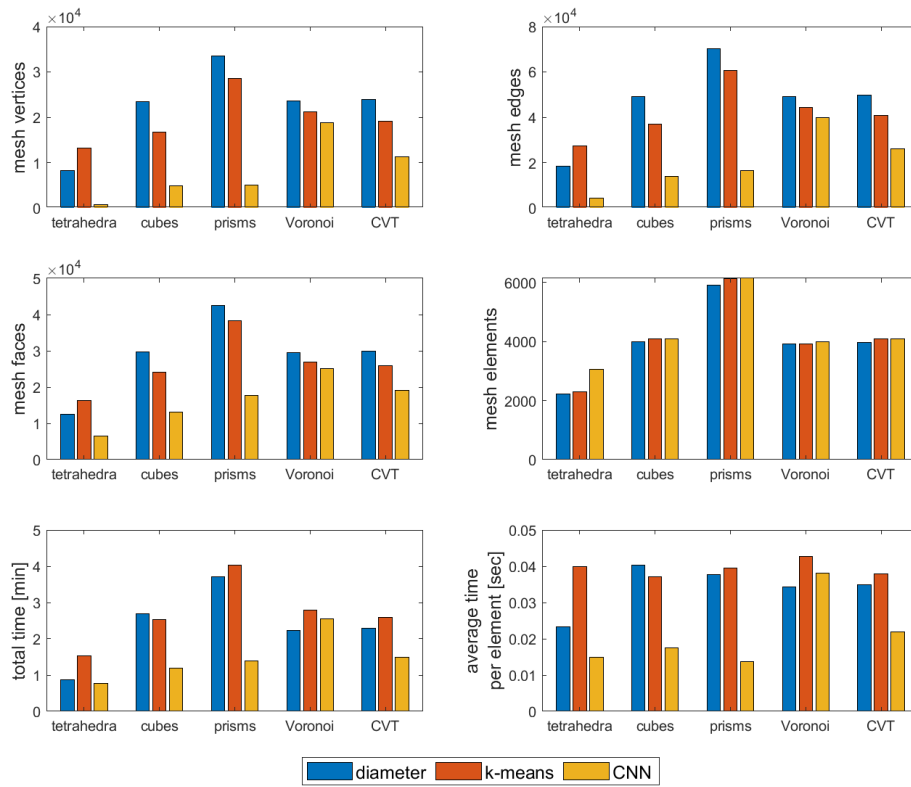


Figure 10: Statistics of computational complexity (number of vertices, edges, faces, elements, total refinement time and average refinement time per element) for the refined grids reported in Figure 8 (second to fourth column), obtained based on employing different refinement strategies (diameter, k-means, CNN).

strategy has the highest computational complexity, which is comparable to the one of the k-means strategy. The diameter strategy is supposed to be the fastest in computing the cutting direction, but because it produces low quality elements with lost of vertices the refinement time becomes comparable with that of the k-means strategy. The CNN-enhanced strategy outperforms both of them in terms of number of geometrical entities and computational time. It is worth noticing that in the CVT the CNN-enhanced strategy achieves higher or comparable quality with respect to the k-means strategy (Figure 9 bottom) in much less time. This is because the CNN is classifying CVT elements as “cubes”, therefore reducing the computational cost. This means that the CNN is generalizing properly on new samples, because CVT elements were not included in the training set.

Summary. The diameter strategy has the simplest implementation but it produces low quality elements, raising the computational cost. The k-means strategy produces the most rounded elements and performs better in unstructured grids. It is robust but has a considerable computational cost. The CNN-enhanced strategy well preserves the mesh structure, producing simple elements of uniform size at a low computational cost. However, it is less effective on unstructured grids.

5 Applications to VEM and PolyDG method

In this section we test the effectiveness of the proposed refinement strategies within polyhedral finite element discretizations. We consider the Virtual Element Method (VEM) [16, 17, 18, 19, 64, 65] and the Polygonal Discontinuous Galerkin (PolyDG) method [9, 10, 11, 12, 13, 14] to solve a standard Laplacian problem in 3D: find $u \in H_0^1(\Omega)$ such that

$$\int_{\Omega} \nabla u \cdot \nabla v = \int_{\Omega} f v \quad \forall v \in H_0^1(\Omega), \quad (1)$$

with $f \in L^2(\Omega)$ a given forcing term. We take $\Omega = (0, 1)^3$ and consider a uniform refinement process in Section 5.1 and an adaptive a priori mesh adaptation procedure in Section 5.2. In both cases, given a grid of the domain Ω we compute numerically the solution of problem (1) using either the VEM or the PolyDG method. We then compute the error, in the VEM case in the discrete H_0^1 -like norm [17], while in the PolyDG case in the DG norm

$$\|v\|_{\text{DG}}^2 = \sum_P \|\nabla v\|_{L^2(P)}^2 + \sum_F \|\gamma^{1/2} \llbracket v \rrbracket\|_{L^2(F)}^2,$$

where P is a polyhedral mesh element and F is a polygonal element face [66, 67]. Here $\gamma = \alpha p^2 / C_{el}$ is the stabilization function, where α is the penalty parameter, p is the polynomial degree used for the approximation, and C_{el} is a function of the shape of the element and is chosen as in [12]. The jump operator $\llbracket \cdot \rrbracket$, which measures the discontinuity of v across elements, is defined as in [66].

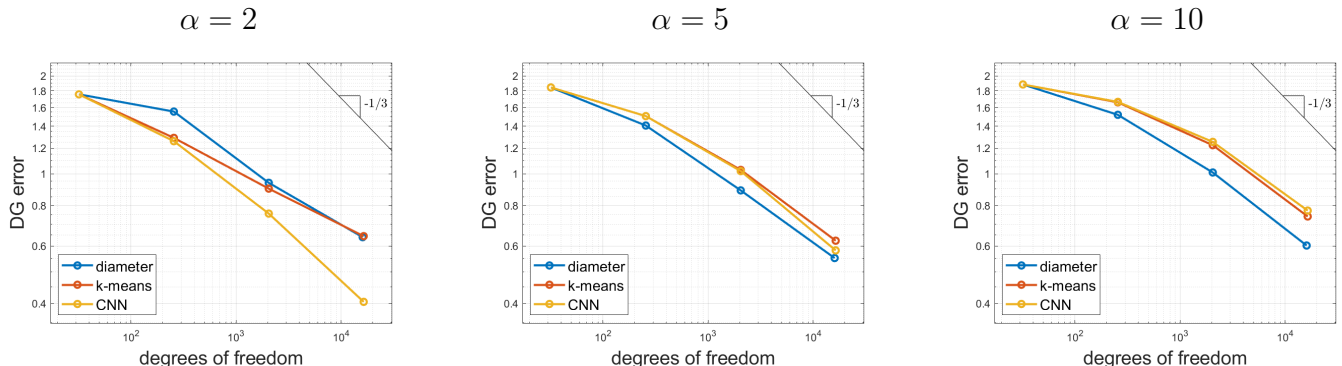


Figure 11: Uniform refinement test case of Section 5.1. Computed errors as a function of the number of degrees of freedom. The PolyDG method is applied to the grid of cubes for different values of the penalty parameter $\alpha = 2, 5, 10$, using polynomials of degree 1. The choice of the most effective refinement strategy (diameter, k-means, CNN) depends on the value of α .

5.1 Uniform refinement

In this Section we consider a uniform refinement process, i.e., at each refinement step each mesh element is refined. The forcing term f in (1) is selected in such a way that the exact solution is given by

$$u(x, y) = \sin(\pi x) \sin(\pi y) \sin(\pi z).$$

The grids obtained after three steps of uniform refinement are those already reported in Figure 8 (second to fourth column).

We now consider the PolyDG method applied to the grid of cubes for different values of the penalty parameter α , using polynomials of degree 1. Larger values of α penalize more discontinuities of the solution, rather than its gradient. In Figure 11 we show the computed errors as a function of the number of degrees of freedom. As we can see, the CNN-enhance strategy is more effective for $\alpha = 2$, while the diameter strategy for $\alpha = 10$. In order to guarantee convergence, α needs to be larger than a threshold that depends on the shape of the elements. For example, $\alpha = 2$ works for the CNN-enhanced strategy, because it preserves the cube structure of the grid, but not for the other strategies. Since computing the threshold is in general too complex, in the following we will use $\alpha = 5$.

In Figure 12 we show the computed errors as a function of the number of degrees of freedom both for the VEM and the PolyDG method, for all grids reported in Figure 8, using “polynomials” of degree 1. When using the PoyDG method, the CNN-enhanced strategy is the most effective in the tetrahedra grid, while the diameter strategy in all other case. To maximize the performance, one may design a CNN-enhanced strategy that classifies polyhedra into two classes only, namely “tetrahedron” and “other”, in order to apply the “classical” strategy in former case and the diameter strategy in the latter. However, considering that the performance of the different strategies are overall comparable and also that

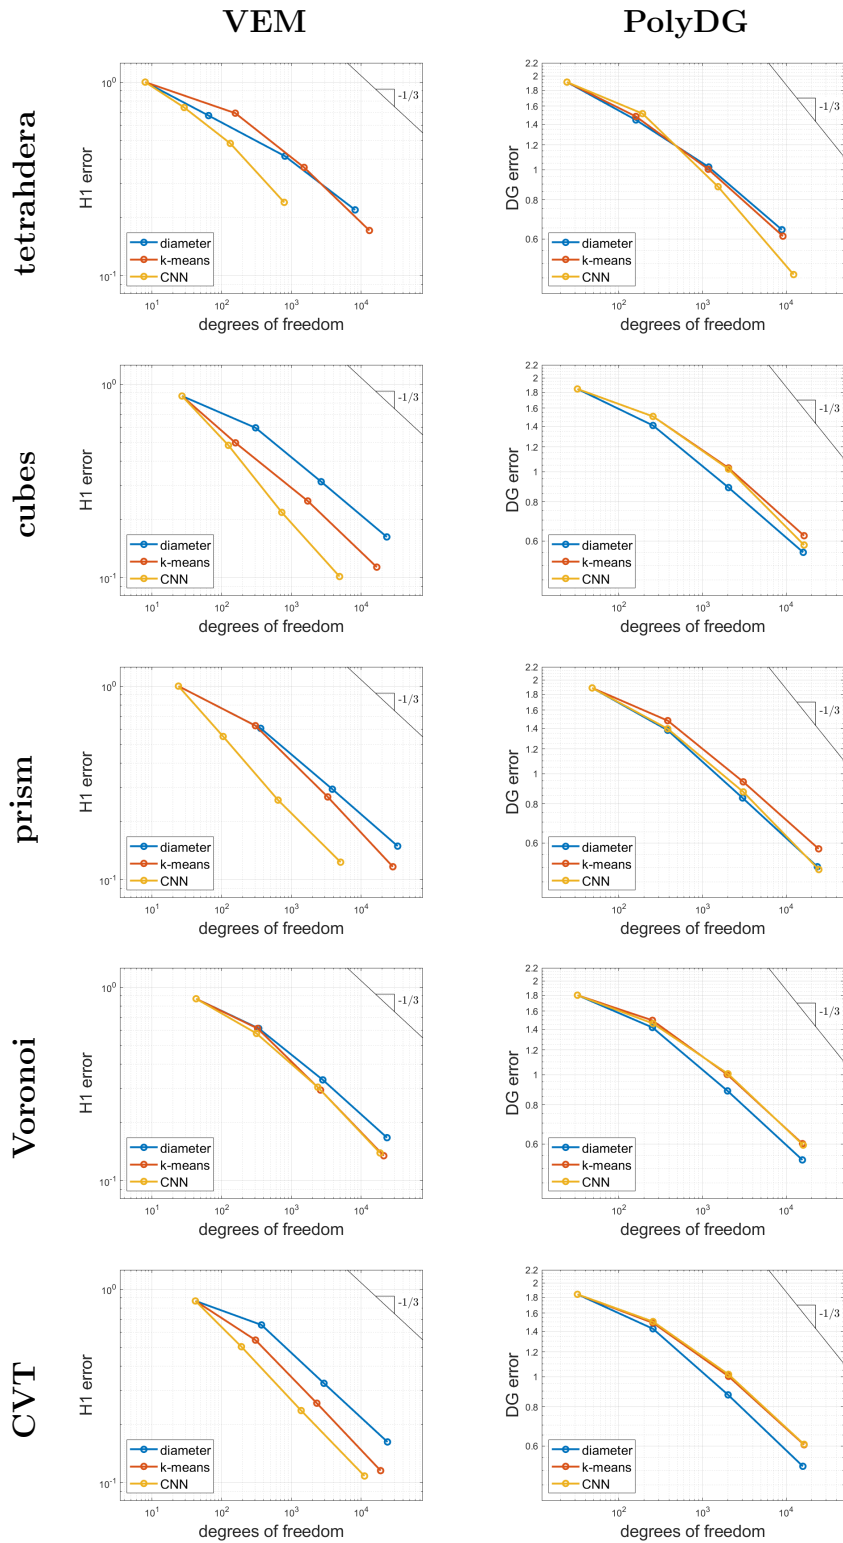


Figure 12: Uniform refinement test case of Section 5.1. Computed errors as a function of the number of degrees of freedom. Each row corresponds to the same initial grid refined uniformly with the proposed refinement strategies (diameter, k-means, CNN), while each column corresponds to a different numerical method (VEM left and PolyDG right).

these results may vary depending on the choice of α , further analysis would be required in order to effectively exploit the use of ML techniques when employing the PolyDG method. On the contrary, in the VEM case the CNN-enhanced strategy significantly outperforms the others, followed by the k-means strategy. At each step the error obtained with the CNN-enhanced refinement is associated with a lower number of degrees of freedom. Indeed, the error lines are shifted to the left. This sensitivity to mesh distortions may be due to the fact that the VEM is a hybrid method, with unknowns on the elements boundary. In Figure 13 we show that the same results hold in the VEM case also for approximation orders 2 and 3.

5.2 Adaptive refinement

In this section we consider an adaptive refinement process. In particular, we proceed as follows:

1. Given a grid of the domain Ω we compute numerically the solution of problem (1) using either the VEM method or the PolyDG.
2. We compute the error with respect to the exact solution, in the VEM case in the discrete H_0^1 -like norm [17], while in the PolyDG case in the DG norm [12]. Notice that we did not employ any a posteriori estimator of the error, since we would like to investigate the effect of the proposed refinement strategies.
3. We refine a fixed fraction r of elements with the highest error. To refine the marked elements we employ one of the proposed strategies (diameter, k-means, CNN).

The forcing term f in (1) is selected in such a way that the exact solution is given by

$$u(x, y) = (1 - e^{-10x})(x - 1) \sin(\pi y) \sin(\pi z),$$

that exhibits a boundary layer along $x = 0$.

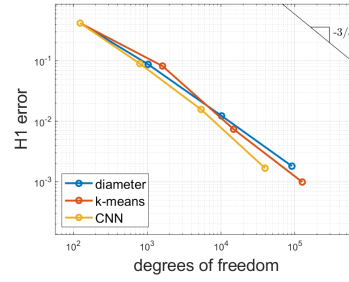
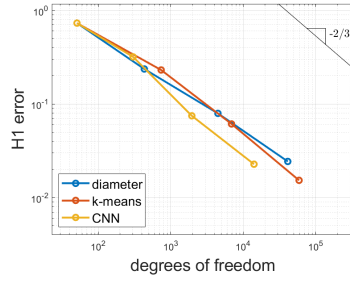
We take as initial grids the first ones of the previous example and we repeat Steps 1-3 for four times using refinement ratio $r = 0.4$. The adapted meshes employing the PolyDG method are shown in Figure 14, similar grids are obtained with VEM. In Figures 15 and 16 we show the computed errors for the PolyDG method of degree 1 and the VEM of orders 1, 2 and 3. Results are analogous to the uniform refinement case: in the VEM case there is a clear advantage in using the CNN-enhanced refinement strategy. Also in this case at each step the error obtained with the CNN-enhanced refinement is associated with a lower number of degrees of freedom: the error lines are shifted to the left. In the PolyDG case all strategies have comparable performance.

VEM

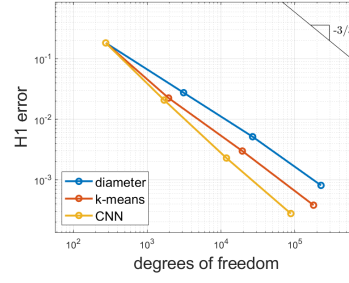
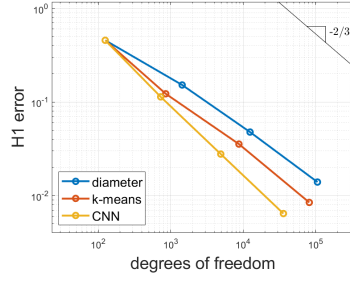
order 2

order 3

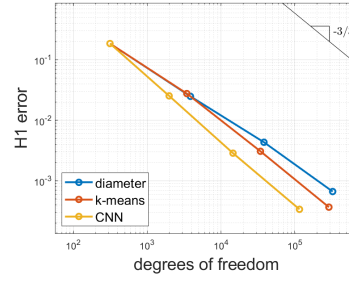
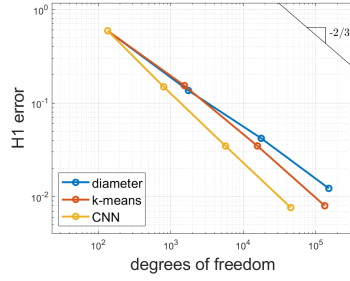
tetrahedra



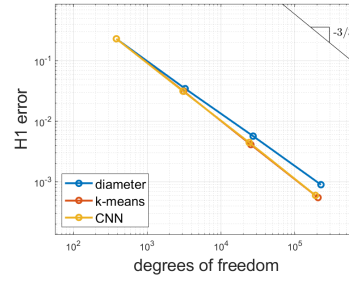
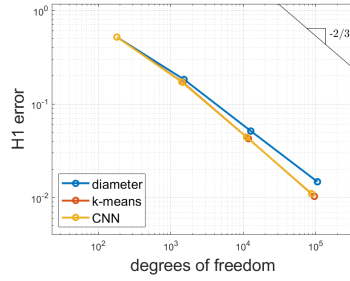
cubes



prism



Voronoi



CVT

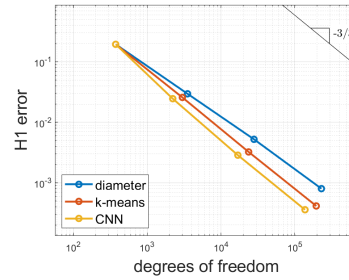
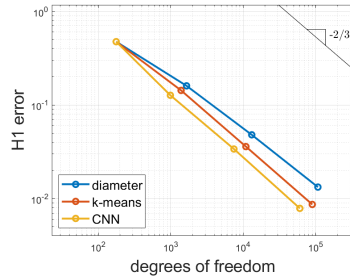


Figure 13: Uniform refinement test case of Section 5.1. Computed errors as a function of the number of degrees of freedom. Each row corresponds to the same initial grid refined uniformly with the proposed refinement strategies (diameter, k-means, CNN), while column corresponds to employing the VEM of order 2 and 3.

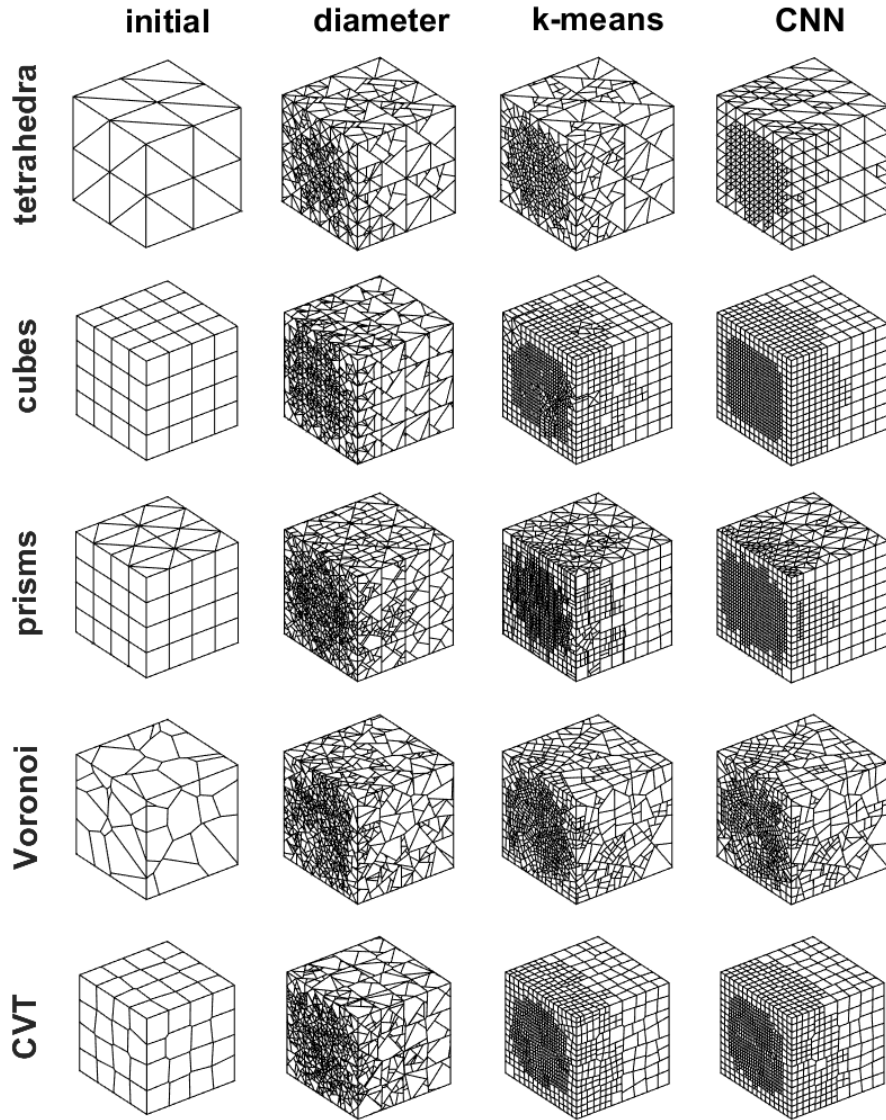


Figure 14: Adaptive refinement test case of Section 5.2. Each row corresponds to the same initial grid (tetrahedra, cubes, prisms, Voronoi, CVT), while each column corresponds to the same refinement strategy (diameter, k-means, CNN). Three steps of adaptive refinement have been performed, with a fixed fraction refinement criterion of $r = 0.4$.

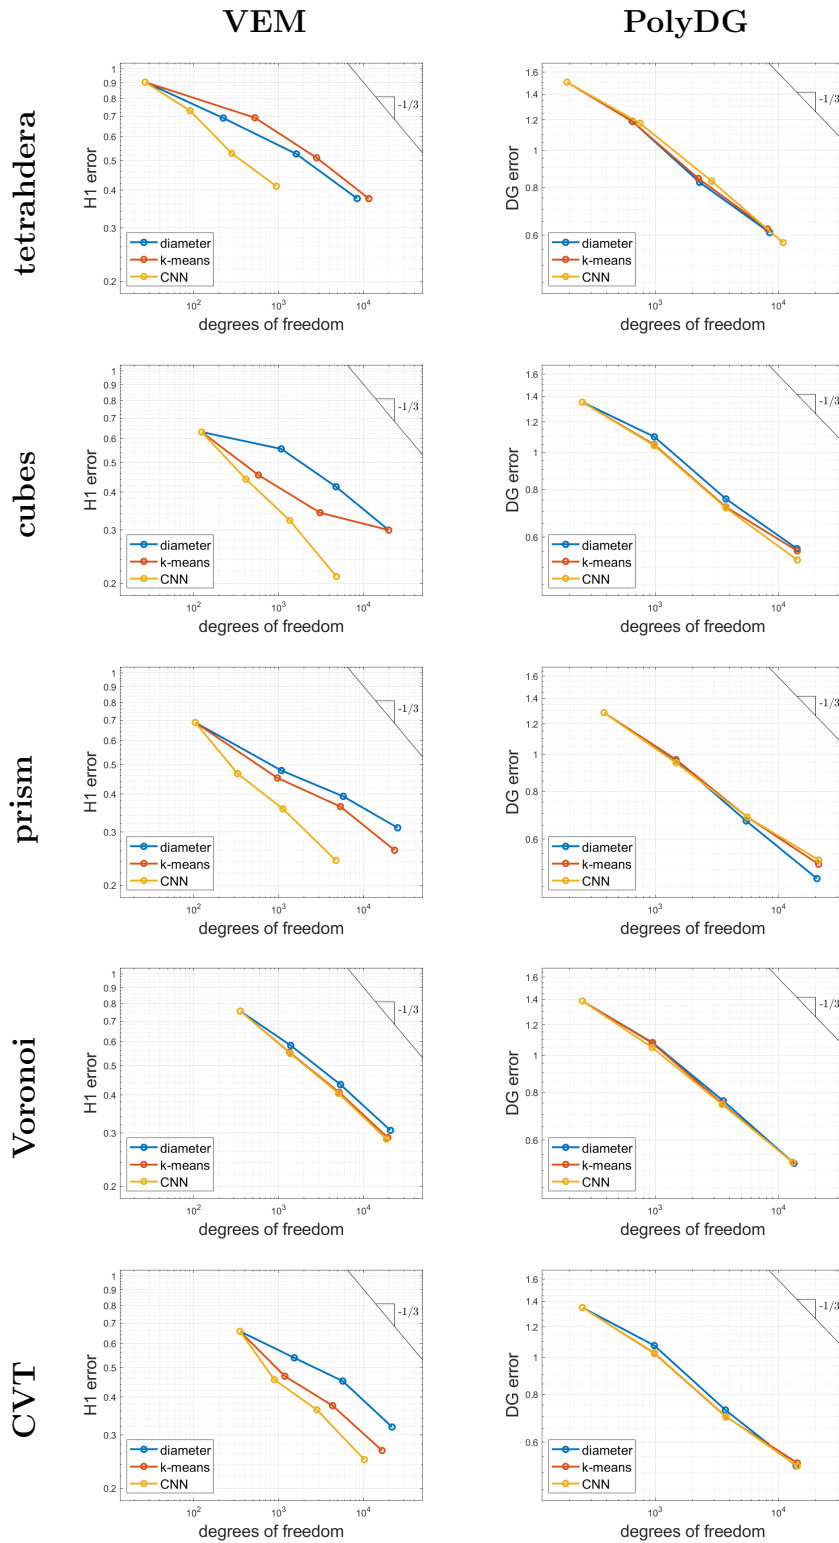


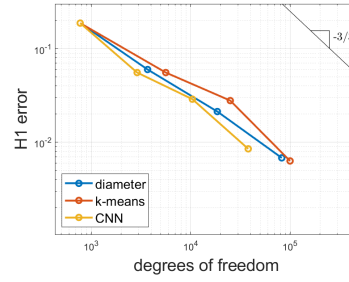
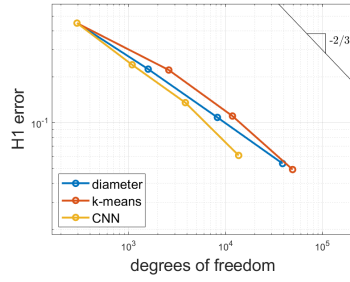
Figure 15: Adaptive refinement test case of Section 5.1. Computed errors as a function of the number of degrees of freedom. Each row corresponds to the same initial grid refined uniformly with the proposed refinement strategies (diameter, k-means, CNN), while each column corresponds to a different numerical method (VEM left and PolyDG right).

VEM

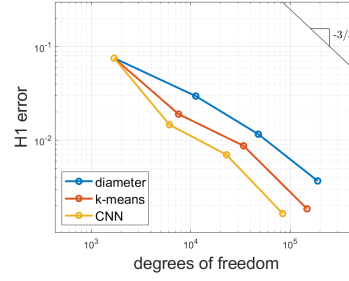
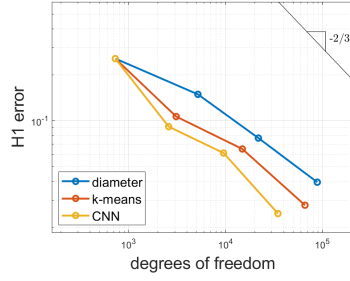
order 2

order 3

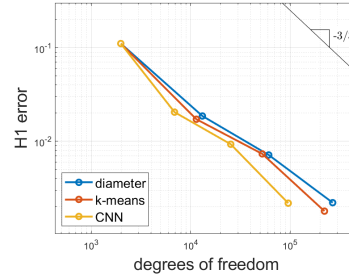
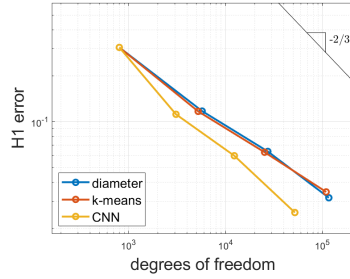
tetrahedra



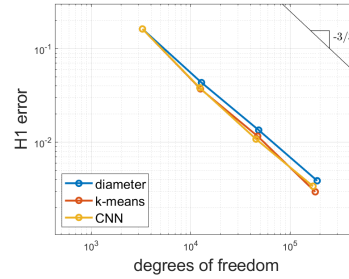
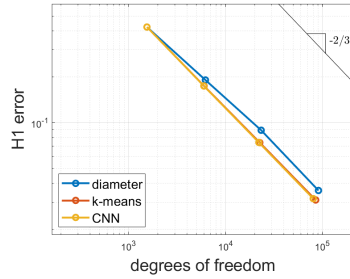
cubes



prism



Voronoi



CVT

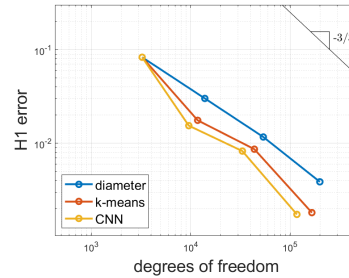
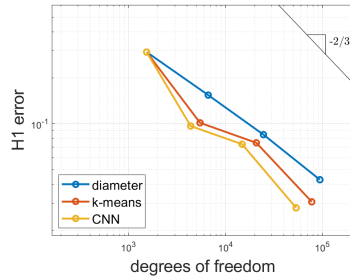


Figure 16: Adaptive refinement test case of Section 5.1. Computed errors as a function of the number of degrees of freedom. Each row corresponds to the same initial grid refined uniformly with the proposed refinement strategies (diameter, k-means, CNN), while column corresponds to employing the VEM of order 2 and 3.

6 Conclusions

In this work, we propose the extension to three dimensions of a paradigm based on ML techniques to enhance existing polygonal grid refinement strategies [42], within polyhedral finite element discretizations of partial differential equations. In particular, the k-means algorithm is used to learn a clustered representation of the element that is used to perform the partition. This strategy is a variation of the well known Centroidal Voronoi Tessellation. It produces shape-regular elements and it is robust when applied on unstructured grids.

Another approach consists in using a CNN as a classifier for polyhedra, that associates to each element the most suitable refinement criterion, including the k-means. This approach has the advantage of being modular: any refinement strategy can be employed, as well as any numerical method can be used to solve the differential problem. Moreover, the CNN has low computational cost once trained, which makes it appealing especially in three dimensions where processing mesh elements is much more expensive than in two dimensions. The use of CNNs allows to exploit strategies explicitly based on the “shape” of the polyhedron, which would not be possible otherwise unless explicit and costly geometrical checks are performed. As a result, the initial structure and quality of the grid is preserved, significantly lowering the computational cost and the number of vertices, edges and faces required to refine the mesh. These ML-enhanced refinement strategies are particularly beneficial for the VEM, which is sensitive to mesh distortions and/or vertex proliferation, as the error is decreasing faster using the same number of degrees of freedom. On the other hand, as expected the PolyDG method is less sensitive as the discretization space is not associated with any geometrical entity.

In terms of future research lines, we plan to use neural networks to drive agglomeration strategies, either by using CNNs to classify the shape of polyhedra or by using graph neural networks to classify the connectivity graph of mesh elements. These will be of paramount importance when designing multilevel solvers.

CRedit authorship contribution statement

P.F. Antonietti: Conceptualization, Funding acquisition, Methodology, Project administration, Resources, Supervision, Writing - review and editing.

F. Dassi: Software, Validation, Visualization, Writing - review and editing.

E. Manuzzi: Conceptualization, Data curation, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

Funding: P.F. Antonietti has been partially supported by the Ministero dell'Università e della Ricerca [PRIN grant numbers 201744KLJL and 20204LN5N5]. P. F. Antonietti, F. Dassi and E. Manuzzi are members of INDAM-GNCS. We thank L. Beirão da Veiga for his support.

References

- [1] J. Hyman, M. Shashkov, and S. Steinberg. “The numerical solution of diffusion problems in strongly heterogeneous non-isotropic materials”. In: *Journal of Computational Physics* 132.1 (1997), pp. 130–148.
- [2] F. Brezzi, K. Lipnikov, and V. Simoncini. “A family of mimetic finite difference methods on polygonal and polyhedral meshes”. In: *Mathematical Models and Methods in Applied Sciences* 15.10 (2005), pp. 1533–1551.
- [3] F. Brezzi, K. Lipnikov, and M. Shashkov. “Convergence of the mimetic finite difference method for diffusion problems on polyhedral meshes”. In: *SIAM Journal on Numerical Analysis* 43.5 (2005), pp. 1872–1896.
- [4] L. Beirao da Veiga, K. Lipnikov, and G. Manzini. *The mimetic finite difference method for elliptic problems*. Vol. 11. Springer, 2014.
- [5] B. Cockburn, B. Dong, and J. Guzmán. “A superconvergent LDG-hybridizable Galerkin method for second-order elliptic problems”. In: *Mathematics of Computation* 77.264 (2008), pp. 1887–1916.
- [6] B. Cockburn, J. Guzmán, and H. Wang. “Superconvergent discontinuous Galerkin methods for second-order elliptic problems”. In: *Mathematics of Computation* 78.265 (2009), pp. 1–24.
- [7] B. Cockburn, J. Gopalakrishnan, and R. Lazarov. “Unified hybridization of discontinuous Galerkin, mixed, and continuous Galerkin methods for second order elliptic problems”. In: *SIAM Journal on Numerical Analysis* 47.2 (2009), pp. 1319–1365.
- [8] B. Cockburn, J. Gopalakrishnan, and F.-J. Sayas. “A projection-based error analysis of HDG methods”. In: *Mathematics of Computation* 79.271 (2010), pp. 1351–1367.
- [9] J. S. Hesthaven and T. Warburton. *Nodal discontinuous Galerkin methods: algorithms, analysis, and applications*. Springer Science & Business Media, 2007.
- [10] F. Bassi et al. “On the flexibility of agglomeration based physical space discontinuous Galerkin discretizations”. In: *Journal of Computational Physics* 231.1 (2012), pp. 45–65.
- [11] P. F. Antonietti, S. Giani, and P. Houston. “hp-version composite discontinuous Galerkin methods for elliptic problems on complicated domains”. In: *SIAM Journal on Scientific Computing* 35.3 (2013), A1417–A1439.
- [12] A. Cangiani, E. H. Georgoulis, and P. Houston. “hp-version discontinuous Galerkin methods on polygonal and polyhedral meshes”. In: *Mathematical Models and Methods in Applied Sciences* 24.10 (2014), pp. 2009–2041.
- [13] P. F. Antonietti et al. “Review of discontinuous Galerkin finite element methods for partial differential equations on complicated domains”. In: *Building bridges: connections and challenges in modern approaches to numerical partial differential equations*. Springer, 2016, pp. 281–310.

- [14] A. Cangiani et al. *hp-Version discontinuous Galerkin methods on polygonal and polyhedral meshes*. Springer, 2017.
- [15] P. F. Antonietti et al. “High-order discontinuous Galerkin methods on polyhedral grids for geophysical applications: seismic wave propagation and fractured reservoir simulations”. In: *Polyhedral Methods in Geosciences* (2021), pp. 159–225.
- [16] L. Beirão da Veiga et al. “Basic principles of virtual element methods”. In: *Mathematical Models and Methods in Applied Sciences* 23.01 (2013), pp. 199–214.
- [17] L. Beirão da Veiga et al. “The hitchhiker’s guide to the virtual element method”. In: *Mathematical models and methods in applied sciences* 24.08 (2014), pp. 1541–1573.
- [18] L. Beirão da Veiga et al. “Virtual element method for general second-order elliptic problems on polygonal meshes”. In: *Mathematical Models and Methods in Applied Sciences* 26.04 (2016), pp. 729–750.
- [19] L. Beirão da Veiga et al. “Mixed virtual element methods for general second order elliptic problems on polygonal meshes”. In: *ESAIM: Mathematical Modelling and Numerical Analysis* 50.3 (2016), pp. 727–747.
- [20] L. Beirão da Veiga et al. “Recent results and perspectives for virtual element methods”. In: *Mathematical Models and Methods in Applied Sciences* (2021), pp. 1–6.
- [21] P. F. Antonietti, L. Beirão da Veiga, and G. Manzini. *The Virtual Element Method and its Applications*. Springer International Publishing, 2022.
- [22] D. A. Di Pietro, A. Ern, and S. Lemaire. “An arbitrary-order and compact-stencil discretization of diffusion on general meshes based on local reconstruction operators”. In: *Computational Methods in Applied Mathematics* 14.4 (2014), pp. 461–472.
- [23] D. A. Di Pietro and A. Ern. “A hybrid high-order locking-free method for linear elasticity on general meshes”. In: *Computer Methods in Applied Mechanics and Engineering* 283 (2015), pp. 1–21.
- [24] D. A. Di Pietro and A. Ern. “Hybrid high-order methods for variable-diffusion problems on general meshes”. In: *Comptes Rendus Mathématique* 353.1 (2015), pp. 31–34.
- [25] D. A. Di Pietro, A. Ern, and S. Lemaire. “A review of hybrid high-order methods: formulations, computational aspects, comparison with other methods”. In: *Building bridges: connections and challenges in modern approaches to numerical partial differential equations*. Springer, 2016, pp. 205–236.
- [26] D. A. Di Pietro and J. Droniou. *The Hybrid High-Order method for polytopal meshes*. Vol. 19. Springer, 2019.
- [27] M. Attene et al. “Benchmark of Polygon Quality Metrics for Polytopal Element Methods”. In: *arXiv preprint arXiv:1906.01627* (2019).

- [28] D. A. Di Pietro, L. Formaggia, R. Masson, et al. “Polyhedral Methods in Geosciences”. In: (2021).
- [29] M.-J. Lai and G. Slavov. “On recursive refinement of convex polygons”. In: *Computer Aided Geometric Design* 45 (2016), pp. 83–90.
- [30] T. Y. Hoshina, I. F. Menezes, and A. Pereira. “A simple adaptive mesh refinement scheme for topology optimization using polygonal meshes”. In: *Journal of the Brazilian Society of Mechanical Sciences and Engineering* 40.7 (2018), p. 348.
- [31] S. Berrone, A. Borio, and A. D’Auria. “Refinement strategies for polygonal meshes applied to adaptive VEM discretization”. In: *Finite Elements in Analysis and Design* 186 (2021), p. 103502.
- [32] T. F. Chan, J. Xu, and L. Zikatanov. “An agglomeration multigrid method for unstructured grids”. In: *Contemporary Mathematics* 218 (1998), pp. 67–81.
- [33] P. F. Antonietti et al. “An agglomeration-based massively parallel non-overlapping additive Schwarz preconditioner for high-order discontinuous Galerkin methods on polytopic grids”. In: *Mathematics of Computation* (2020).
- [34] M. Raissi, P. Perdikaris, and G. E. Karniadakis. “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations”. In: *Journal of Computational Physics* 378 (2019), pp. 686–707.
- [35] M. Raissi and G. E. Karniadakis. “Hidden physics models: Machine learning of nonlinear partial differential equations”. In: *Journal of Computational Physics* 357 (2018), pp. 125–141.
- [36] F. Regazzoni, L. Dedè, and A. Quarteroni. “Machine learning for fast and reliable solution of time-dependent differential equations”. In: *Journal of Computational Physics* 397 (2019), p. 108852.
- [37] F. Regazzoni, L. Dedè, and A. Quarteroni. “Machine learning of multi-scale active force generation models for the efficient simulation of cardiac electromechanics”. In: *Computer Methods in Applied Mechanics and Engineering* 370 (2020), p. 113268.
- [38] J. S. Hesthaven and S. Ubbiali. “Non-intrusive reduced order modeling of nonlinear problems using neural networks”. In: *Journal of Computational Physics* 363 (2018), pp. 55–78.
- [39] D. Ray and J. S. Hesthaven. “An artificial neural network as a troubled-cell indicator”. In: *Journal of Computational Physics* 367 (2018), pp. 166–191.
- [40] P. F. Antonietti, M. Caldana, and L. Dedè. “Accelerating Algebraic Multi-grid Methods via Artificial Neural Networks”. In: *arXiv preprint arXiv:2111.01629* (2021).

- [41] F. Regazzoni et al. “A machine learning method for real-time numerical simulations of cardiac electromechanics”. In: *arXiv preprint arXiv:2110.13212* (2021).
- [42] P. Antonietti and E. Manuzzi. “Refinement of polygonal grids using Convolutional Neural Networks with applications to polygonal Discontinuous Galerkin and Virtual Element methods”. In: *Journal of Computational Physics* 452 (2022), p. 110900. ISSN: 0021-9991.
- [43] J. A. Hartigan and M. A. Wong. “Algorithm AS 136: A k-means clustering algorithm”. In: *Journal of the royal statistical society. series c (applied statistics)* 28.1 (1979), pp. 100–108.
- [44] A. Likas, N. Vlassis, and J. J. Verbeek. “The global k-means clustering algorithm”. In: *Pattern recognition* 36.2 (2003), pp. 451–461.
- [45] Q. Du, V. Faber, and M. Gunzburger. “Centroidal Voronoi tessellations: Applications and algorithms”. In: *SIAM review* 41.4 (1999), pp. 637–676.
- [46] Y. Liu et al. “On centroidal Voronoi tessellation—energy smoothness and fast computation”. In: *ACM Transactions on Graphics (ToG)* 28.4 (2009), pp. 1–17.
- [47] Y. LeCun, Y. Bengio, and G. Hinton. “Deep learning”. In: *nature* 521.7553 (2015), pp. 436–444.
- [48] S. Albawi, T. A. Mohammed, and S. Al-Zawi. “Understanding of a convolutional neural network”. In: *2017 International Conference on Engineering and Technology (ICET)*. Ieee. 2017, pp. 1–6.
- [49] S. C. Brenner and L.-Y. Sung. “Virtual element methods on meshes with small edges or faces”. In: *Mathematical Models and Methods in Applied Sciences* 28.07 (2018), pp. 1291–1336.
- [50] L. Beirão da Veiga, C. Lovadina, and A. Russo. “Stability analysis for the virtual element method”. In: *Mathematical Models and Methods in Applied Sciences* 27.13 (2017), pp. 2557–2594.
- [51] J. Droniou and L. Yemm. “Robust Hybrid High-Order method on polytopal meshes with small faces”. In: *arXiv preprint arXiv:2102.06414* (2021).
- [52] L. Mu, X. Wang, and Y. Wang. “Shape regularity conditions for polygonal/polyhedral meshes, exemplified in a discontinuous Galerkin discretization”. In: *Numerical Methods for Partial Differential Equations* 31.1 (2015), pp. 308–325.
- [53] S. C. Johnson. “Hierarchical clustering schemes”. In: *Psychometrika* 32.3 (1967), pp. 241–254.
- [54] F. Murtagh and P. Contreras. “Algorithms for hierarchical clustering: an overview”. In: *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 2.1 (2012), pp. 86–97.
- [55] H. Ritter, T. Martinetz, K. Schulden, et al. *Neural computation and self-organizing maps: an introduction*. Addison-Wesley Reading, MA, 1992.

- [56] J. Kangas, T. Kohonen, and J. Laaksonen. “Variants of self-organizing maps”. In: *IEEE transactions on neural networks* 1.1 (1990), pp. 93–99.
- [57] J. Lane, B. Magedson, and M. Rarick. “An efficient point in polyhedron algorithm”. In: *Computer vision, graphics, and image processing* 26.1 (1984), pp. 118–125.
- [58] N. Le Hoang, T. K. Dang, et al. “A farthest first traversal based sampling algorithm for k-clustering”. In: *2020 14th International Conference on Ubiquitous Information Management and Communication (IMCOM)*. IEEE, 2020, pp. 1–6.
- [59] C. B. Barber, D. P. Dobkin, and H. Huhdanpaa. “The quickhull algorithm for convex hulls”. In: *ACM Transactions on Mathematical Software (TOMS)* 22.4 (1996), pp. 469–483.
- [60] I. L. Dryden and K. V. Mardia. *Statistical shape analysis: with applications in R*. Vol. 995. John Wiley & Sons, 2016.
- [61] C. M. Bishop. *Pattern recognition and machine learning*. springer, 2006.
- [62] P. C. Petersen. “Neural network theory”. In: *University of Vienna* (2020).
- [63] D. P. Kingma and J. Ba. “Adam: A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980* (2014).
- [64] L. B. Da Veiga, F. Dassi, and A. Russo. “High-order virtual element method on polyhedral meshes”. In: *Computers & Mathematics with Applications* 74.5 (2017), pp. 1110–1122.
- [65] F. Dassi and L. Mascotto. “Exploring high-order three dimensional virtual elements: bases and stabilizations”. In: *Computers & Mathematics with Applications* 75.9 (2018), pp. 3379–3401.
- [66] D. N. Arnold et al. “Unified analysis of discontinuous Galerkin methods for elliptic problems”. In: *SIAM Journal on Numerical Analysis* 39.5 (2002), pp. 1749–1779.
- [67] B. Cockburn, G. E. Karniadakis, and C.-W. Shu. *Discontinuous Galerkin methods: theory, computation and applications*. Vol. 11. Springer Science & Business Media, 2012.

MOX Technical Reports, last issues

Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

- 11/2022** Sampaoli, S.; Agosti, A.; Pozzi, G.; Ciarletta, P.
A toy model of misfolded protein aggregation and neural damage propagation in neurodegenerative diseases
- 10/2022** Fresca, S.; Manzoni, A.
Real-time simulation of parameter-dependent fluid flows through deep learning-based reduced order models
- 09/2022** Corti, M.; Zingaro, A.; Dede', L.; Quarteroni, A.
Impact of Atrial Fibrillation on Left Atrium Haemodynamics: A Computational Fluid Dynamics Study
- 07/2022** Sinigaglia, C.; Quadrelli, D.E.; Manzoni, A.; Braghin, F.
Fast active thermal cloaking through PDE-constrained optimization and reduced-order modeling
- 06/2022** Pozzi, G.; Grammatica, B.; Chaabane, L.; Catucci, M.; Mondino, A.; Zunino, P.; Ciarletta, P.
T cell therapy against cancer: a predictive diffuse-interface mathematical model informed by pre-clinical studies
- 08/2022** Gobat, G.; Opreni, A.; Fresca, S.; Manzoni, A.; Frangi, A.
Reduced order modeling of nonlinear microstructures through Proper Orthogonal Decomposition
- 05/2022** Aspri, A.; Beretta, E.; Cavaterra, C.; Rocca, E.; Verani, M.
Identification of cavities and inclusions in linear elasticity with a phase-field approach
- 04/2022** Africa, P.C.; Piersanti, R.; Fedele, M.; Dede', L.; Quarteroni, A.
lifex - heart module: a high-performance simulator for the cardiac function
- 01/2022** Gavazzoni, M.; Ferro, N.; Perotto, S.; Foletti, S.
Multi-physics inverse homogenization for the design of innovative cellular materials: application to thermo-mechanical problems
- 03/2022** Giacomini, M.; Perotto, S.
Anisotropic mesh adaptation for region-based segmentation accounting for image spatial information