



Reinforcement learning for block decomposition of planar CAD models

Benjamin C. DiPrete¹ · Rao Garimella² · Cristina Garcia Cardona³ · Navamita Ray⁴

Received: 20 April 2023 / Accepted: 25 December 2023 / Published online: 14 February 2024
This is a U.S. Government work and not under copyright protection in the US; foreign copyright protection may apply 2024

Abstract

The problem of hexahedral mesh generation of general CAD models has vexed researchers for over 3 decades and analysts often spend more than 50% of the design-analysis cycle time decomposing complex models into simpler blocks meshable by existing techniques. The decomposed blocks are required for generating good quality meshes (tilings of quadrilaterals or hexahedra) suitable for numerical simulations of physical systems governed by conservation laws. We present a novel AI-assisted method for decomposing (segmenting) planar CAD (computer-aided design) models into well shaped rectangular blocks. Even though the simple examples presented here can also be meshed using many conventional methods, we believe this work is proof-of-principle of a AI-based decomposition method that can eventually be generalized to complex 2D and 3D CAD models. Our method uses reinforcement learning to train an agent to perform a series of optimal cuts on the CAD model that result in a good quality block decomposition. We show that the agent quickly learns an effective strategy for picking the location and direction of the cuts and maximizing its rewards. This paper is the first successful demonstration of an agent autonomously learning how to perform this block decomposition task effectively, thereby holding the promise of a viable method to automate this challenging process for more complex cases.

Keywords Quadrilateral and Hexahedral Mesh generation · Block decomposition · Reinforcement learning

1 Introduction

Many numerical methods compute approximate solutions over a mesh of topologically simpler elements (tetrahedra, hexahedra) representing the computational domain. In highly non-linear problems (e.g. fluid dynamics with shocks), hexahedra are preferred, or even required, over tetrahedra because of their superior accuracy and directional control of the solution [1]. In spite of 30+ years of research, however, there are no reliable algorithms that can automatically generate hexahedral meshes for general CAD models [2, 3]. Contrast this with tetrahedral meshing which has long been automatic at scale for realistic industrial problems [4, 5].

In an early paper, Thompson [6] proposed a multi-block grid generation method to generate hexahedral meshes for geometric model naturally composed of 6-sided blocks that are topologically cubical but with general geometry. Each block in the domain is meshed by mapping a structured mesh of cube to the general geometry using transfinite mapping [7] while also ensuring that the meshes are continuous across block boundaries. Another early paper by White

✉ Rao Garimella
rao@lanl.gov

Benjamin C. DiPrete
bendiprete@gmail.com

Cristina Garcia Cardona
cgarcia@lanl.gov

Navamita Ray
nray@lanl.gov


¹ College of Computing, Georgia Institute of Technology, Atlanta, GA, USA

² Applied Computational Physics Division, XCP-4, Los Alamos National Laboratory, Los Alamos 87544, New Mexico, USA

³ Computer, Computational and Statistical Sciences Division, CCS-3, Los Alamos National Laboratory, Los Alamos 87544, New Mexico, USA

⁴ Computer, Computational and Statistical Sciences Division, CCS-7, Los Alamos National Laboratory, Los Alamos 87544, New Mexico, USA

Fig. 1 **a** CAD model that cannot be automatically meshed entirely with hexahedra **b** model is subdivided into 6-sided blocks using cuts **c** each block is meshed by mapping a regular hexahedral mesh of a unit cube onto the block



et al. [8] proposed the reverse process by which realistic 3D geometric models were virtually decomposed into 6-sided blocks and then tiled with hexahedra (Fig. 1). This process, called Block Decomposition, is guided by human intuition and acquired domain expertise that readily “sees” how to subdivide a model for a particular application. Attempts to automate this process have not proven generalizable to arbitrary shapes [9, 10].¹

1.1 Previous work

There have been sustained efforts [5] to develop automatic algorithms to generate hexahedral meshes for complex geometric models since 3D Finite Element Methods became popular. An elementary method (ca. 1980) called mapped meshing uses transfinite interpolation to map the structured mesh of a canonical cube to topologically equivalent but geometrically different domains [11]. For roughly tubular geometric models, an algorithm called multi-sweeping [12, 13] extrudes a quadrilateral mesh on one set of faces to form stacks of elements that reach an opposite set.

The Block Decomposition [8] method targeted here generalizes these techniques by decomposing complex geometric models into parts that are amenable to mapped meshing or multi-sweeping. Block decomposition is favored by seasoned analysts for its superior control of mesh quality and directionality despite the fact that it must be done manually. While there have been significant efforts to devise automatic decomposition algorithms of complex models based on the model characteristics directly [8–10, 14–16] or on alternate model representations like the medial-axis transform [17, 18], most methods have remained experimental or work on a limited class of problems.


¹ In general we will use the word model or shape to implicitly mean a geometric model and explicitly point out when we use it to mean a machine learned approximation of reality.

In recent years, there has been a sharp uptick in research into using artificial intelligence (AI) or machine learning (ML) with deep neural networks (NN) for solving meshing related problems. Much of the work has focused on using AI/ML for generating or tweaking 2D triangular meshes with point densities suited for a particular PDE (partial differential equation) solution bypassing mesh adaptation using *a posteriori* error estimation [19–25]. Pan et al. [26, 27] describe an actual ML-based quadrilateral mesh generation method. A more recent paper by Tong et al. [28] uses a combination of supervised learning and reinforcement learning to assist the advancing front method for generating high quality quadrilateral meshes without the need for complex checks like front intersection. There are some older papers claiming to use “knowledge-based methods” to generate meshes [29, 30], recognize model features [31, 32], or even decompose geometric models [14, 33] but none of them used ML as we know it. Recent papers on CAD and ML have focused mainly on Shape Matching [34–36] and to a lesser extent on CAD model generation [37, 38] and cleanup [39, 40].

1.2 Our approach

This article presents a proof-of-principle demonstration of a novel AI-assisted method for decomposing complex geometric models into blocks by applying it to planar shapes with straight, axis-aligned edges. Our approach uses reinforcement learning (RL) [41] to let an agent learn a good sequence of steps to take in order to cut the input model into meshable blocks. In RL, an agent learns by taking actions in an environment based on the state of the environment. Each action moves the environment into a new state and grants a reward to the agent. With a targeted balance of exploration vs exploitation, the agent learns a policy that maximizes the cumulative reward over a sequence of actions. RL closely mimics how human analysts learn to decompose complex shapes into blocks and

Fig. 2 Recursive slicing of the model. The left figure shows the original model and a vertex at which the agent is poised to act along with the two cuts it can make. The middle figure shows the two shapes from the first action and the choices for the next action. The right figure shows the shapes arising from the second cut



in recent years, RL, combined with deep neural networks (DNN), has matched or surpassed human-level skill in several fields [42]. It is worth noting that this study is different from the use of reinforcement learning for image segmentation in medicine [43] or in video processing [44, 45] or segmentation of 3D point clouds [46].

There are many challenges in applying reinforcement learning to the problem of block decomposition of complex geometric models. Unlike common scenarios like learning to play a game or navigate a warehouse where the environment is fixed, our environment is dynamically changing as we make cuts. Thus a naively formulated global observation set (i.e. the data about the evolving geometric configuration that we can feed to the agent) will vary in size as the episode progresses making it unsuitable for traditional neural networks. The agent itself has multiple types of decisions to make - where to perform a modification and what type of modification to make (full cut, partial cut, etc.). Additionally, the parameters of the modification are continuous (for example, the angle of a cut) and the agent must learn a distribution of expected rewards over the continuous parameter space. Finally, the task of the agent is not to master the decomposition of one particular geometric model - rather the ultimate goal is to learn a generalizable policy that can be applied to new configurations.

To tackle this problem, we devise an RL agent to process an input geometric model that is planar with straight, axis-aligned edges. The agent recursively subdivides it into simpler parts using axis-aligned cuts. The environment is a custom setup that can read a geometric model and answer queries about it (e.g. how many vertices, how many edges connected to a vertex, the angle formed by two edges at a vertex). The agent also uses the capabilities of the geometric modeler to make modifications to the shape - in this particular study, the modification is slicing the geometric model into two or more pieces from a model vertex. The quality of the resulting parts (reduced complexity, low aspect ratio) determines the reward the agent receives. An episode ends when the input is decomposed into all quadrilateral blocks. In the results section, we demonstrate that our RL agent

quickly learns which cuts to make and where to make them to maximize its rewards.

While the method is currently demonstrated on simple problems that may be solved using procedural algorithms such as the art gallery algorithm [47], the purpose of this paper is not necessarily to demonstrate superior quality or performance in the decomposition of these simple shapes. Rather it is to introduce an AI framework that encapsulates most of the principles required to apply it to more complex 2D and 3D shapes and demonstrate that we can effectively tackle diverse planar configurations without needing to adapt the formulation on a case-by-case basis. We believe this is the first time such a reinforcement learning approach has been used to tackle the problem of block decomposition.

2 Methodology

We have developed a customized RL framework that learns how to effectively decompose geometric models into blocks by exploring the effect of different geometric model modifications. While most components of our RL framework are set up for general problems in 2D and 3D, this study is limited to decomposing planar shapes with straight, axis-aligned edges. The CAD model is described using a full-featured 3D geometric modeler called OpenCascade [48] but for the purposes of this discussion, it can be considered to be one or more planar shapes, each of which is described by a sequence of model vertices and model edges. During each step of the training phase, the agent picks a vertex of the geometric model, observes the state and makes a geometric modification. Currently, the only geometric modification the agent can make is a full cut, i.e., slice the geometric model into two or more parts using an infinite line (See Fig. 2). While we use an RL technique that allows for a continuous action space (e.g. cuts originating at any location and angled arbitrarily), we restrict the cuts in this study to only originate from a model vertex and be aligned with the X- or Y-axis. Since the geometric model evolves as the agent makes cuts, the size of a global observation set for the full model, e.g.

the list of vertices, also changes and cannot be used directly as input to a traditional neural network. Therefore, following the idea of Pan [27], we have designed a fixed size local observation set at each model vertex to feed to the neural networks in the RL framework. The iterative application of this sequence of steps - select vertex, construct local observations, make a cut, evaluate the quality - allows the agent to learn to block decompose the geometric model. In order to learn a policy to efficiently perform such a decomposition, the agent is trained via feedback from the environment: cuts that produce a *good partition*, e.g. resulting in quadrilateral blocks with good aspect ratios are *rewarded*, while cuts that produce a *bad partition*, e.g. high aspect ratios in its decomposed parts, high variance in the areas of its decomposed parts or cuts that do not affect the model (cutting along a side) are *penalized*. The policy learned in this way can then be applied to perform block decomposition of other planar axis-aligned shapes.

2.1 Soft actor-critic-based RL architecture

Our framework uses the soft actor-critic (SAC) reinforcement learning algorithm introduced in [49]. The SAC method provides a sample efficient (i.e. moderate data collection demands) and stable, model-free,² deep RL algorithm for continuous state and action spaces. While it may be argued that this problem might be tackled with a deep Q-network, the reason for using a SAC-type algorithm is to build a framework that can be generalized to more complex 2D and 3D models that require arbitrarily angled or partial cuts from any boundary location.

There are three main components in the SAC algorithm:

1. An actor-critic architecture with separate policy and value function networks,
2. An off-policy formulation that enables reuse of previously collected data for efficiency, and
3. Entropy maximization to enable stability and exploration.

The implementation of the soft actor-critic architecture includes three separate networks: an *actor* network, a *critic* network and a *value function* network that are optimized jointly during training. As discussed by [49], this not only provides flexibility to handle large continuous domains, but can also help to stabilize training.

² Here we are talking of a model of a general environment for RL, not a geometric model.

2.1.1 Actor network

The actor network outputs a probability distribution over the action space \mathcal{A} and is also in charge of executing actions. In our case, it is implemented as a traditional neural network that receives as input a local observation (described below). Its output determines the probability for each of the two directions allowed for cuts from a given vertex: along the X-axis or Y-axis. Note that training uses a stochastic actor, where the selection of a cutting direction is made randomly weighted by the estimated probabilities, while, during deployment, the actor behaves deterministically selecting the action with the maximum estimated probability. The stochasticity is useful to maximize the entropy of the actor network and encourage exploration of the environment in the training phase.

2.1.2 Critic network

The critic network qualifies how good the allowed actions are for a given state. It is similar to a Q-network in Deep-Q learning [42] in that it learns to approximate the Q-value of actions in a given state, i.e. it learns to approximate the reward for a given *state-action pair*,³ along with all future rewards along the expected trajectory. In our case, it is also implemented as a traditional neural network that receives as input a local observation and determines the Q-value (quality) of X-axis and Y-axis cuts.

2.1.3 Local observation

The actor and critic networks are represented as traditional neural networks that expect a fixed input structure and, thus, are not able to handle the varying size and complexity of the evolving environment (i.e. the changing collection of vertices and edges as the geometric model is sliced repeatedly). Hence, we construct a special fixed structure to capture important *local* shape information *observed* at a chosen model vertex. The features included in this structure are:

- Vectors to the two neighboring vertices
- Type of interior angle formed by the two vectors (acute, right, obtuse, reentrant, etc)
- Vector to the centroid of the shape being processed
- Aspect ratio of the shape being processed

A schematic of the local observation features can be found in Fig. 3. As explained later, the complexity of observations

³ In reinforcement learning terminology, an action-state pair refers to the selection of a particular action from all possible actions in a particular state of the environment.

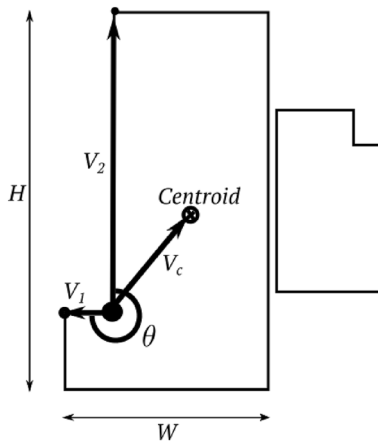



Fig. 3 Features included in the local observation: Vectors to neighboring vertices (V_1, V_2) and vector to centroid of the shape (V_c), angle of the vertex corner (θ), aspect ratio of full shape (H/W)

at model vertices in our study remains fixed because the two parts resulting from a cut are treated as independent parts for the next cut - thus every model vertex remains connected to two adjacent vertices.

2.1.4 Value network

The value network qualifies how good a particular state is. In other words, this network learns to approximate the expected reward and future rewards the actor will receive in a given state. In our case, this network allows the actor to choose the next vertex to perform a cut. Thus, it is more appropriate to regard this network as being able to approximate the expected reward and future rewards the actor will receive for making a cut from a specific vertex. For efficiently capturing all the relevant vertex-level information for the full model, this network must be able to handle the varying collection of vertices produced during shape decomposition. Hence, we implement this network as a graph neural network (GNN), specifically as a SplineCNN network [50]. The network receives as input a triangular mesh of the planar model. We

Fig. 4 Example planar shape and corresponding triangulation input to the value network. In the triangulation input, red mesh vertices lie on model vertices (vertex type 2), green mesh vertices lie on model edges (vertex type 1), blue mesh vertices lie in the interior (vertex type 0)



can control the resolution of this triangular mesh, usually preferring coarse meshes to avoid excessive computational burden. We tag the mesh vertices as being coincident with model vertices, lying on a model edge or lying in the interior of the model as shown in Fig. 4b. Furthermore, notice that the GNN not only allows us to work with a changing number of vertices, it also enables the incorporation of spatial geometric information of the current decomposition state, information that would be much more difficult to encode using a traditional NN.

Although the value network produces an output at every mesh node, only the outputs at the model vertices (i.e. red points in Fig. 4b) are considered. As stated above, the output value of the value network at a model vertex is an approximation to the expected reward and future rewards if a cut is made at that vertex. With this structure in place, we can chose a vertex to perform a cut at every step of an episode. Mimicking the stochastic actor concept, the set of values produced by the value network on the model vertices is used during training as probability weights and the vertex to perform a cut is randomly selected using these weights with the goal of encouraging exploration. In contrast, the selection is deterministic during deployment and the vertex with the highest output of the value network is selected to perform a cut.

2.2 Off-policy formulation

The SAC algorithm uses off-policy actor-critic training, combined with a stochastic actor as described before, which results in a more stable and scalable algorithm. Such a strategy allows it to reuse past experience to train the models and increases the sample efficiency. It is implemented by storing a distribution \mathcal{D} of previously sampled states, actions and rewards, and using it as a replay buffer during training. We follow this approach during training which alternates between collecting experience from the environment by applying the current policy, and updating the networks via

stochastic gradients computed from batches sampled from the replay buffer.

2.3 Entropy Maximization

Unlike the regular actor-critic framework, SAC rewards entropy in its actions by optimizing policies to maximize both the expected return and the expected entropy of the policy. This encourages exploration of the environment and makes the algorithm more robust and capable of general learning, rather than just memorization. The maximum entropy policies are also robust to estimation errors and improve exploration by allowing the acquisition of diverse behaviors.

2.4 Reward Function

The reward function is a critical component of the RL framework and contributes to the effectiveness with which the agent carries out the task at hand. In our case, we devise a reward function to

- Encourage creating quadrilateral parts
- Discourage cuts that do not affect the geometric model (e.g. cutting along a side)
- Discourage high variance in the areas of its decomposed parts
- Discourage high aspect ratios in its decomposed parts

Once the geometric model is fully decomposed into blocks, the agent gets a bonus reward and the episode concludes. The exact form of the rewards used for this study are given in the results section.

2.5 Training Phase

The training phase is composed of a collection of episodes, each *episode* consisting of all the steps needed for decomposing a given geometric model. During a training episode, the agent uses the value network output to select a vertex to cut, and the actor network output to select the particular action to take. Both of these are done stochastically to ensure a higher level of exploration during training.

The steps listed below are iterated during a training episode

1. Triangulate the shape being processed
2. Run the value network on the triangulation to generate weights at mesh vertices
3. Stochastically select a model vertex based on value network outputs
4. Compile a local observation at the vertex

5. Stochastically choose a direction for a cut at the vertex based on actor network probability outputs
6. Split the geometric model into two or more parts along the chosen direction
7. Compute the new state and reward
8. Store sampled states, actions and rewards in the replay buffer
9. Update parameters for every network following the gradient step
10. Pick another non-quadrilateral part from the geometric model decomposition and repeat from step 1

Geometric models are loaded repeatedly from the training set, one per episode. A set number of episodes is run during training. The training of all the networks uses the Adam optimization algorithm. The functions optimized in each case are the same as in the SAC original work. There is, however, a slight difference in the value network: when calculating the loss, the network only propagates loss for the node that was chosen to make a cut from.


Note that a cut goes fully through the shape and splits it into two or more parts (see Fig. 2). Instead of keeping the model as a collection of generated parts, we treat each part as a new shape to explore. Thus at each step we split the model, set aside quadrilateral parts, and put the remaining parts in a processing queue. This approach sacrifices the full model view, but makes it simpler and more robust since the agent does not encounter a local state of ever increasing complexity and there is no need to accumulate the knowledge of how the parts build up. An additional benefit of this approach is that each new part generated is a training data sample for the agent.

2.6 Deploying the Trained Framework

After the framework is trained, the combination of value network and actor network constitute the learned policy. The decomposition of new geometric model proceeds as follows (with similarities to the training phase):

1. Triangulate the shape being processed
2. Run the value network on the triangulation to generate weights at mesh vertices
3. Deterministically select a model vertex with highest value output by the value network
4. Compile a local observation at the vertex
5. Deterministically choose the cut direction with the highest probability as predicted by the actor network
6. Split the geometric model into two or more parts along the chosen direction
7. Compute the new state and reward
8. Pick another non-quadrilateral part from the geometric model decomposition and repeat from step 1

Fig. 5 Samples from the (a) training data set containing 37 models and (b) test data set containing 12 models



Crucially, at the end of the decomposition, all the shapes are merged back together while retaining the boundaries between them. Thus vertices that appear on the boundary of one block are also reflected in the boundary of adjacent blocks. The merged model is then meshed using well-known procedures. In our case, we import the parts into the CUBIT geometric modeling and meshing package [51], use its imprint-and-merge functionality to recreate a single geometric model (with internal cuts) and apply a mapped meshing algorithm.

3 Numerical Experiments

3.1 Data Sets

Our training and testing data set includes 49 planar shapes with straight, axis-aligned edges. These were generated using our python script that invokes the CUBIT package [51] to randomly generate and combine 2 to 10 rectangles. The training and test data sets consist of 37 models and 12 independent models respectively (Figs. 5a, 5b).

3.2 Network Architecture

All networks are implemented using PyTorch [52] and PyTorch Geometric [53]. The architectures used are described next.

3.2.1 Actor and Critic Networks

These networks are traditional feed-forward NN composed of 4 fully connected layers, with 256, 128, 64 and 2 neurons, respectively (the last of these layers is the output layer). We use rectified linear unit (ReLU) activation functions after each of these layers, except for the last layer in the critic


network⁴ that uses a linear activation function. The input dimension is 9, corresponding to the size of the local observation: 2 dimensional (2D) vector for each of the 2 neighboring vertices, 2D vector to centroid, 1 value for angle at vertex and 2 components to represent the aspect ratio. The networks have 2 outputs which correspond to the dimension of the action space (i.e. 2 cut directions: X-axis or Y-axis).

3.2.2 Value Network

This network is a GNN. It contains 1 SplineCNN layer, followed by 7 residual blocks and 1 final SplineCNN output layer. Each residual block is composed of 2 SplineCNN layers. There are batch normalization layers after all the SplineCNN layers except the output layer. We use exponential linear unit (ELU) activation functions except in the output layer. Every SplineCNN layer has a kernel size of 5, meaning the 2D B-spline function for the continuous kernel has 25 defining points, with 5 points on each axis. The number of nodes in the graph is arbitrary and depends on the triangulation of the shape. Each node in the graph input layer has 3 features because each node has one-hot encoded vector features: (1, 0, 0) represents interior point, (0, 1, 0) represents boundary point and (0, 0, 1) represents model vertex. Each node in the graph output layer has 1 feature corresponding to the value function for that node, but only nodes corresponding to model vertices are taken into account. The first SplineCNN layer has 64 features. The residual blocks have 128, 256, 128, 64, 32, 16 and 8 features, respectively. Note that if the number of features does not change through a residual block, the input features to the residual block are simply summed with the output features. However, if the number of features does change through a block, the skip connection contains 1 SplineCNN layer, with as many features as the features in the block. All our residual blocks change the number of features.


⁴ We actually use two equally parameterized critic networks as in the original SAC work, see [49] for details.

Fig. 6 Reward convergence obtained by the training



(a) Rewards obtained by the agent on the training dataset as it trained. (b) Rewards obtained by the agent on the test dataset as it trained.

Fig. 7 The block decomposition (middle) returned by the agent for the test shape shown on the left and its mesh from CUBIT (right)



3.2.3 Reward Function

Assume that a splitting action on a shape results in N new shapes, with N_q of them being quadrilaterals. Let the areas of the shapes be A_i , $i = 1, N$, and aspect ratios R_i , $i = 1, N$ (where the aspect ratio of a shape is defined as the ratio of the longest side to the shortest side of its bounding box). Also, let the average area of all the shapes be \bar{A} .

The reward \mathcal{R} is defined as

$$3 \left[\left(\frac{N}{\sum_i R_i^2} \right)^{\frac{1}{2}} - \frac{(\sum_i (A_i - \bar{A})^2)^{\frac{1}{2}}}{\sum_i A_i} - 1 \right] + 10 \frac{N_q}{N} - 5\delta_{1N} \quad (1)$$

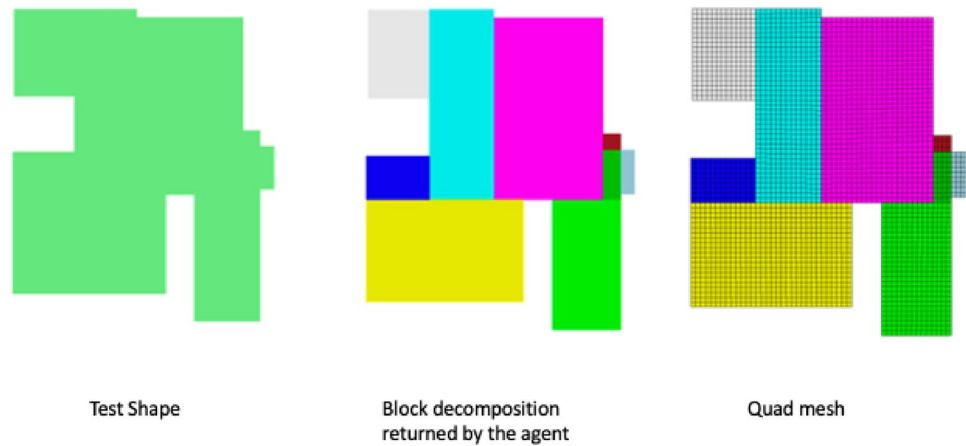
Note that minimum possible aspect ratio is 1 and therefore the leading term (reciprocal of the root mean square of aspect ratios) takes a maximum value of 1 when all shapes are squares. The second term which measures variance in the areas of the shapes takes a minimum value of 0 when all the areas are equal. The third term is a maximum if the action results in all quads ($N_q = N$). The fourth term serves as a penalty for actions that result in no new shapes ($N = 1$). Thus the maximum reward is obtained when the action cuts the shape into squares of equal area.

3.3 Testing and Reward Convergence

As the model learns using the training set, the RL framework’s learning is periodically checked against the test set. In a testing episode, the vertex at which to act and the action to take are chosen deterministically to maximize the reward - a vertex with the highest output from the value network is chosen, and the action with the highest probability from the actor network is applied at that vertex.

Figure 6a shows a moving average of rewards (over 10 episodes) obtained by the RL framework during the training phase. Figure 6b shows the convergence of a moving average of rewards during the periodic testing episodes. We observe that after only around 1500 episodes of training (an hour or so of training time) the model learned to obtain consistently high rewards on its training set, but also on the test set of shapes it has never trained on. The oscillations in the reward plot of the training set indicate that the agent is continuing to favor exploration rather than exploitation. The good reward convergence seen on the test set implies that the agent is steadily gathering generalizable knowledge about the decomposition problem for this category of shapes.

Fig. 8 The block decomposition (middle) returned by the agent for the test shape shown on the left and its mesh from CUBIT (right)



3.4 Decomposition Examples

Finally, in Figs. 7 and 8, we present two examples of block decompositions obtained for test shapes (i.e. shapes that the agent never trained on). It showcases the learned knowledge of the agent after it was trained. The block decompositions were then meshed using CUBIT to generate quadrilateral meshes of the decomposed shape.

4 Conclusions

We have demonstrated a novel reinforcement learning-based AI method to decompose input CAD shapes into well shaped blocks that can be meshed for numerical simulations. The results show that an agent using the SAC reinforcement learning framework can learn a block decomposition policy that generalizes to new planar, axis-aligned shapes.

While this proof-of-principle demonstration is restricted to simple 2D shapes and elementary geometric model modifications, it contains most of the principles required to generalize it to more complex shapes in 2D and 3D. The environment is based on geometric modelers which regularly handle complex 3D shapes with curved boundaries. The agent's actions are modeled on the types of operations a human agent decomposing a shape will execute using a geometric modeler (e.g. planar model cuts). The use of Soft-Actor-Critic framework allows for continuous actions (e.g. cuts at an angle) in the future. Similarly, the rewards are based on the quality evaluation of the blocks used by meshing algorithms and analysts. The issue of variability in the starting environment and the dynamic evolution of the environment are already addressed in this simple problem using a graph-based value neural network. Thus, we can reasonably surmise that the method can eventually be generalized to address the real problem of decomposing 3D shapes thereby alleviating one of the long standing problems in meshing.

5 Future Work

In the future, we will expand this research to tackle more complex 2D and 3D shapes. We will extend this method to non-axis aligned 2D shapes by first cutting along edges and eventually at arbitrary angles. Expanding to more complex curved geometric models will require expansion of the types of actions to include partial cuts or some other templated subdivision (like making a square internal boundary inside a circular part). The reward function definitions may also have to be refined further. Expanding the method to 3D requires tetrahedral meshes for the value network, an expanded set of observations, generalized reward functions and more types of geometric modifications.

Funding This work was supported by the U.S. Department of Energy for Los Alamos National Laboratory (LANL) under contract 89233218CNA000001. This publication is approved for release as LANL report number LA-UR-23-29571. The authors have no relevant financial or non-financial interests to declare regarding the topic of research in this article.

Data availability Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Wang E, Nelson T, Rauch R (2004) Back to elements - tetrahedra vs hexahedra. In: Proceedings of the 2004 International Ansys Conference. <https://www.ansys.com/-/media/ansys/corporate/resourcelibrary/conference-paper/2004-int-ansys-conf-9.pdf>
- Pietroni N, Campen M, Sheffer A, Cherchi G, Bommers D, Gao X, Scateni R, Ledoux F, Remacle J-F, Livesu M (2022) Hex-mesh generation and processing: a survey. *ACM Transactions on Graphics* 42(2):1–44
- Sarrate, J., Ruiz-Gironés, E., Roca, X.: Unstructured and semi-structured hexahedral mesh generation methods. *Computational Technology Reviews* 10 (2017)
- Shephard MS, Georges MK (1991) Automatic three-dimensional mesh generation by the finite octree technique. *International Journal for Numerical Methods in Engineering* 32(4):709–749
- Owen S (2016) An Introduction to Automatic Mesh Generation Algorithms. Short Course Notes International Meshing Roundtable, Washington, D.C. . <https://www.osti.gov/servlets/purl/1394107>
- Thompson JF (1987) A general three-dimensional elliptic grid generation system on a composite block structure. *Computer Methods in Applied Mechanics and Engineering* 64:377–411
- William G, Hall C (1973) Construction of curvilinear coordinate systems and their application to mesh generation. *International Journal of Numerical Methods in Engineering* 7(4):461–477
- White, D.R., Mingwu, L., Benzley, S.E., Sjaardema, G.D.: Automated hexahedral mesh generation by virtual decomposition. In: Proceedings of the 4th International Meshing Roundtable, pp. 165–176 (1995)
- White DR, Saigal S, Owen SJ (2004) Ccsweep: automatic decomposition of multi-sweep volumes. *Engineering with Computers* 20:222–236
- Wang R, Shen C, Chen J, Wu H, Gao S (2017) Sheet operation based block decomposition of solid models for hex meshing. *Computer-Aided Design* 85:123–137
- Gordon WJ, Thiel LC (1982) Transfinite mappings and their application to grid generation. *Applied Mathematics and Computation* 10–11:171–233
- Mingwu, L., Benzley, S.E., Sjaardema, G., Tautges, T.: A multiple source and target sweeping method for generating all hexahedral finite element meshes. In: Proceedings of 5th International Meshing Roundtable, pp. 165–176 (1996)
- Shepherd, J.F., Mitchell, S.A., Knupp, P.M., White, D.R.: Methods for multisweep automation. In: Proceedings of the 9th International Meshing Roundtable, pp. 77–87 (2000)
- Lu, Y., Gadhd, R., Tautges, T.J.: Volume decomposition and feature recognition for hexahedral mesh generation. In: 8th International Meshing Roundtable, Lake Tahoe, CA, pp. 269–280 (1999)
- White DR, Tautges TJ (2000) Automatic scheme selection for toolkit hex meshing. *International Journal of Methods in Engineering* 49(1):127–144
- Lu JH-C, Quadros WR, Shimada K (2017) Evaluation of user-guided semi-automatic decomposition tool for hexahedral mesh generation. *Journal of Computational Design and Engineering* 4(4):330–338
- Blum, H.: A transformation for extracting new descriptors of shape. In: Models for the Perception of Speech and Visual Forms, pp. 362–380 (1967)
- Price MA, Armstrong CG, Sabin MA (1995) Hexahedral mesh generation by medial surface subdivision: part i. solids with convex edges. *International Journal of Numerical Methods in Engineering* 38(19):3335–3359
- Zhang, Z., Wang, Y., Jimack, P.K., Wang, H.: Meshingnet: A new mesh generation method based on deep learning. *Computational Science – ICCS 2020. Lecture Notes in Computer Science* 12139 (2020)
- Huang, K., Krügener, M., Brown, A., Menhorn, F., Bungartz, H.-J., Hartmann, D.: Machine Learning-based optimal mesh generation in computational fluid dynamics. *arXiv preprint arXiv:2102.12923* (2021)
- Dielen, A., Lim, I., Lyon, M., Kobbelt, L.: Computing direction fields for quad mesh generation. In: Eurographics Symposium on Geometry Processing, vol. 40 (2021)
- Pak, D., Liu, M., Kim, T., Liang, L., McKay, R., Sun, W., Duncan, J.: Distortion energy for deep learning-based volumetric finite element mesh generation for aortic valves. *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021. Lecture Notes in Computer Science* 12906, 485–494 (2021)
- Bohn J, Feischl M (2021) Recurrent neural networks as optimal mesh refinement strategies. *Computers & Mathematics with Applications* 97:61–76
- Yang, J., Dzanic, T., Petersen, B., Kudo, J., Mittal, K., Tomov, V., Camier, J.-S., Zhao, T., Zha, H., Kolev, T., et al.: Reinforcement learning for adaptive mesh refinement. In: International Conference on Artificial Intelligence and Statistics, pp. 5997–6014 (2023)
- Wu T, Liu X, An W, Huang Z, Lyu H (2022) A mesh optimization method using machine learning technique and variational mesh adaptation. *Chinese Journal of Aeronautics* 35(3):27–41
- Pan J, Huang J, Wang Y, Cheng G, Zeng Y (2021) A self-learning finite element extraction system based on reinforcement learning. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing* 35:180–208
- Pan, J., Huang, J., Cheng, G., Zeng, Y.: Reinforcement learning for automatic quadrilateral mesh generation: a soft actor-critic approach. *arXiv report* (2022). <https://arxiv.org/abs/2203.11203>
- Tong H, Qian K, Halilaj E, Zhang YZ (2023) SRL-assisted AFM: Generating planar unstructured quadrilateral meshes with supervised and reinforcement learning-assisted advancing front method 72:102109
- Manevitz LM, Yousef M, Givoli D (2002) Finite-element mesh generation using self-organizing neural networks. *Computer-Aided Civil and Infrastructure Engineering* 4(12):233–250
- Çinar, A., Arslan, A.: Neural networks based mesh generation method in 2-d. *Lecture Notes in Computer Science* 2510 (Eur-Asia-ICT 2002: Information and Communication Technology), 395–401 (2002)
- Kim YS (1992) Recognition of form features using convex decomposition. *Computer-Aided Design* 24(9):461–476
- Wu H, Gao S, Wang R, Chen J (2018) Fuzzy clustering based pseudo-swept volume decomposition for hexahedral meshing. *Computer-Aided Design* 96:42–58
- Takata O et al (1999) A knowledge-based mesh generation system for forging simulation. *Applied Intelligence* 11(2):149–168
- Qin, F., Li, L., Gao, S., et. al.: A deep learning approach to the classification of 3d cad models. *Journal of Zhejiang Univ. - Science C* 15, 91–106 (2014)
- Bronstein, M.M., et.al.: Geometric deep learning: Going beyond euclidean data. *IEEE Signal Processing Magazine* 34(4) (2017)
- Boussuge, F., Tierney, C.M., Robinson, T.T., Armstrong, C.G.: Application of tensor factorisation to analyse similarities in cad assembly models. In: Proceedings of the 28th International Meshing Roundtable (2019)
- Wu, R., X., C., , Zheng, C.: DeepCAD: A deep generative network for computer-aided design models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (2021)
- Lin, C., Fan, T., Wang, W., Nießner, M.: Modeling 3d shapes by reinforcement learning. In: *Computer Vision – ECCV 2020*.

- Lecture Notes in Computer Science(), vol. 12355 (2020). https://doi.org/10.1007/978-3-030-58607-2_32
39. Danglade F, Pernot J-P, Véron P (2014) On the use of machine learning to defeature cad models for simulation. *Computer-Aided Design and Applications* 11(3):358–368
 40. Owen, S., Shead, T.M., Martin, S.: CAD defeaturing using machine learning. In: *Proceedings of the International Meshing Roundtable* (2019)
 41. Sutton RS, Barto AG (2018) *Reinforcement Learning: An Introduction*. A Bradford Book, Cambridge, MA, USA
 42. Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G, Petersen S, Beattie C, Sadik A, Antonoglou I, King H, Kumaran D, Wierstra D, Legg S, Hassabis D (2015) Human-level control through deep reinforcement learning. *Nature* 518(7540):529–533. <https://doi.org/10.1038/nature14236>
 43. Sahba, F., Tizhoosh, H.R., Salama, M.M.A.: A reinforcement learning framework for medical image segmentation. In: *The 2006 IEEE International Joint Conference on Neural Network Proceedings*, pp. 511–517 (2006). doi: 10.1109/IJCNN.2006.246725
 44. Han, J., Yang, L., Zhang, D., Chang, X., Liang, X.: Reinforcement cutting-agent learning for video object segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2018)
 45. Wang, Y., Dong, M., Shen, J., Wu, Y., Cheng, S., Pantic, M.: Dynamic face video segmentation via reinforcement learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2020)
 46. Tiator, M., Geiger, C., Grimm, P.: Point cloud segmentation with deep reinforcement learning. In: *24th European Conference on Artificial Intelligence 2020*, pp. 2768–2775 (2020)
 47. Mukherjee N (2014) An art gallery approach to submap meshing. *Procedia Engineering* 82:313–324
 48. OpenCascade: OpenCascade Technology, 7.7.0dev. OpenCascade.com, Moulinaux, France (2022). OpenCascade.com. <https://dev.opencascade.org/doc/overview/html/>
 49. Haarnoja, T., Zhou, A., Abbeel, P., Levine, S.: Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *CoRR abs/1801.01290* (2018) 1801.01290
 50. Fey, M., Lenssen, J.E., Weichert, F., Müller, H.: SplineCNN: Fast geometric deep learning with continuous B-spline kernels. *CoRR abs/1711.08920* (2017) 1711.08920
 51. CUBIT: The Cubit Geometry and Mesh Generation Toolkit. Sandia National Laboratories, Albuquerque, NM, USA (2022). Sandia National Laboratories. https://cubit.sandia.gov/files/cubit/16.06/help_manual/WebHelp/cubithelp.htm
 52. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pp. 8024–8035 (2019)
 53. Fey, M., Lenssen, J.E.: Fast graph representation learning with PyTorch Geometric. In: *ICLR Workshop on Representation Learning on Graphs and Manifolds* (2019)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.