A coupled Lloyd–Poisson framework for metric-based Voronoi mesh adaptation

By A. Stock

A metric-based Voronoi mesh adaptation framework is developed to achieve spatial resolution control consistent with a prescribed target metric. The method combines Lloyd smoothing with a Poisson-based correction step that displaces the Voronoi generators along a potential field derived from the local metric error. This coupling ensures simultaneous convergence of both cell quality and metric conformity. The Voronoi formulation naturally accommodates topology changes through regeneration, avoiding mesh stretching and tangling.

1. Introduction

The accuracy, stability and computational cost of a numerical flow simulation are intrinsically governed by the underlying mesh resolution. In the context of high-fidelity large-eddy simulation (LES), the mesh defines the smallest resolved spatial scale and thus sets the effective cutoff length of the discretized equations. In regions characterized by steep gradients of velocity, temperature or species mass fraction, such as boundary layers, shear interfaces, shocks or flame fronts, a fine resolution is required to correctly capture local dynamics and minimize discretization errors. Conversely, coarse elements can be employed in slowly varying regions to reduce the number of degrees of freedom and the overall computational burden. The optimal mesh is therefore one that minimizes numerical error for a prescribed computational cost, which is equivalent to enforcing a spatially varying resolution that adapts to the flow physics.

A widely adopted mathematical framework for expressing such spatially varying resolution is the metric field, introduced by George *et al.* (2018). The metric tensor $\mathcal{M}(x) \in \mathbb{R}^{d\times d}$ defines at each point an idealized unit sphere in the metric space, such that the length of any edge e satisfies

$$\ell_M(e) = \sqrt{e^T\,\mathcal{M}(x)\,e} \approx 1. \tag{1.1}$$

The eigenstructure of $\mathcal{M}$ prescribes both the target element size via the eigenvalues and the preferred stretching directions via the eigenvectors. Metric-based mesh adaptation thus consists of generating a mesh whose element edge lengths are close to the unit in this metric sense. In the current work, an isotropic metric is considered. As a consequence all diagonal terms of $\mathcal{M}(x)$ are equal, and nondiagonal terms are null, thus effectively reducing the metric tensor to a scalar field, $\mathcal{M}(x) \in \mathbb{R}$.

Numerous approaches have been proposed for constructing suitable metric fields in LES, which can broadly be classified as either feature-based or error-estimator-based. Feature-based indicators target dynamically active regions identified from resolved flow quantities of interest, for example, gradients of velocity or density, flame or shock sensors and interfaces. Metrics derived from subgrid-scale quantities such as turbulence or turbulent combustion models (Bose *et al.* 2011; Agostinelli *et al.* 2021; Stock & Moureau 2024)

act as hybrid indicators: They target unresolved features while also quantifying local modeling error. Similarly, Hessian-based metrics (Habashi *et al.* 2000; Formaggia *et al.* 2004; Loseille & Alauzet 2011; Bénard *et al.* 2016) originate from interpolation-error analysis but can also be interpreted as feature detectors when normalized by a physical scale, as they highlight zones of strong curvature in velocity, pressure or scalar fields. Finally, purely error-estimator-based techniques such as adjoint or goal-oriented adaptation (Hartmann *et al.* 2011; Offermans *et al.* 2023) derive the metric from sensitivity or residual information, ensuring optimal accuracy for target functionals. In practice, adaptive frameworks often blend these criteria to obtain robust metrics that capture both physical and numerical sources of error.
While the definition of a suitable metric is well established, obtaining a discrete mesh that conforms to it remains challenging. Classical adaptive remeshing frameworks typically rely on tetrahedral meshes and iteratively apply local topological operators such as edge splitting, collapsing, swapping and vertex displacements until metric conformity is achieved (Dapogny *et al.* 2014). Although robust, such algorithms are computationally expensive and can struggle to maintain consistency across parallel domain interfaces sharing vertices or edges between processors. Voronoi tessellations (Lloyd 1982) constitute an appealing alternative for unstructured mesh generation and adaptation, providing a continuous and topology-preserving framework for controlling cell size and distribution.

Section 2 introduces the principles of Voronoi mesh generation and examines how both cell quality and metric conformity can be enforced through various adaptation strategies. Section 3 presents the proposed methodology, which achieves convergence toward a prescribed metric via a Poisson-based correction coupled with topology regeneration. Section 4 demonstrates the effectiveness of the approach on a two-dimensional test case, confirming the method's ability to reach the target metric while maintaining good mesh quality.

2. Voronoi mesh generation

2.1. *Voronoi diagram*

A Voronoi diagram is a canonical partitioning of a domain $\Omega \subset \mathbb{R}^d$ induced by a discrete set of generator points $X = \{x_i\}_{i=1}^N \subset \Omega$. The Voronoi cell associated with generator x_i is defined as

$$V_i(X) = \big\{\, x \in \Omega \;\big|\; \|x - x_i\| \leq \|x - x_j\| \;\; \forall j \neq i \big\}. \tag{2.1}$$

Each point $x \in \Omega$ thus belongs to the region corresponding to its nearest generator. The Voronoi tessellation is dual to the Delaunay triangulation, where generators connected by a shared Voronoi face form Delaunay edges.

2.2. *Converging mesh quality*

Beyond spatial resolution, the quality of mesh elements plays a critical role in the stability and accuracy of computational fluid dynamics (CFD) solvers. Lloyd iterations provide an effective means of relaxing highly skewed or stretched cells toward a more regular configuration by updating each generator to the centroid of its associated cell

$$x_i = \frac{\int_{V_i} x \, dx}{\int_{V_i} dx}. \tag{2.2}$$

Repeated iterations converge the mesh toward a stationary state where generators and centroids are merged, known as centroidal voronoi tessellation (CVT). This effectively minimizes the energy functional

$$E(X) = \sum_{i=1}^{N} \int_{V_i} \|x - x_i\|^2 \, dx. \tag{2.3}$$

While this process improves mesh regularity and smoothness, it also promotes uniform cell sizes. Therefore, an additional mechanism is required to enforce convergence to a prescribed metric field.

2.3. *Converging mesh metric*

Convergence to a given metric can be achieved in numerous ways. The main methods are outlined here.

2.3.1. *Generator insertion and deletion*

Voronoi mesh refinement can be done by performing generator insertion and deletion to best satisfy the metric (Labelle & Shewchuk 2003; Mao 2015). This method works best for generating uniform metric layers with metric jumps of a factor of 2, with Lloyd smoothing the transitions.

2.3.2. *Weighted Lloyd*

A weighted Lloyd iteration (Du *et al.* 1999) can be defined as

$$x_i = \frac{\int_{V_i} x \, \rho(x) \, dx}{\int_{V_i} \rho(x) \, dx}, \tag{2.4}$$

where $\rho(x)$ is a prescribed spatial density field. This causes the generators to be shifted up the density gradients until a sufficient metric gradient is present between the high- and low-density regions. This results in smaller cells in high-density regions, which can be leveraged for mesh adaptation (Praud *et al.* 2025). The limitation of this approach is that reaching a specific metric distribution requires iterative adjustment of ρ, with a full Lloyd convergence performed between successive ρ updates.

2.3.3. *Laguerre Diagram*

To incorporate non-uniform resolution, the Voronoi distance can be generalized using additive weights, w,

$$V_i(X) = \left\{ x \in \Omega \mid \|x - x_i\|^2 - w_i \leq \|x - x_j\|^2 - w_j, \ \forall j \neq i \right\}. \tag{2.5}$$

The resulting cells define a Laguerre (or power) diagram. The weights w_i allow control over cell volumes independently of their geometric location, enabling better adaptation to a prescribed metric field.

2.3.4. *Variational mesh adaptation and MMPDE*

Variational mesh adaptation (VMA) (Huang *et al.* 1994) performs a smooth remapping of the computational domain to minimize an energy functional associated with the metric error. This remapping preserves the mesh topology, preventing cell tangling and the formation of negative cell volumes. It can be interpreted as a local stretching, compression and rotation of the mesh to conform to a prescribed metric field. The resulting mesh is anisotropic, with element shapes and orientations aligned with the gradients of the metric.

The moving mesh partial differential equation (MMPDE) method provides a dynamic alternative, performing multiple intermediate remapping steps that gradually relax the mesh toward the optimal configuration. The main limitations of both VMA and MMPDE arise in scenarios involving complex or highly non-uniform metric targets, which are difficult to satisfy through pure mesh deformation.

2.3.5. *Method limitations*

When generating meshes that conform to a prescribed metric field, two mechanisms are generally required. First, point insertion and deletion operations adjust the total number of generators to match the integral count dictated by the metric distribution, producing an initial mesh that roughly satisfies the target density. Second, a Lagrangian displacement step refines the generator positions to achieve local conformity. Weighted Lloyd iterations or Laguerre diagrams enable such displacements, but each generator motion is constrained within its current Voronoi cell, limiting the possible displacement amplitude per iteration. As a result, mesh adaptation proceeds through small, locally driven updates; each generator moves according to the shape of its own cell and the local density gradient, influencing only its immediate neighbors in the next iteration. Consequently, information propagates diffusively across the domain at a rate of one cell layer per iteration, leading to slow global convergence on large meshes. Moreover, the adjustment of the weight or density field governing the displacement is subtle and typically requires repeated convergence cycles.

VMA and MMPDE frameworks overcome this locality by providing a global, continuous mesh deformation that accelerates convergence toward the target metric. However, their fixed-connectivity formulations inherently produce anisotropic and often distorted elements, with the degree of anisotropy emerging as a byproduct of the deformation rather than an explicitly controlled parameter. For many CFD applications, such uncontrolled anisotropy can degrade solver stability and accuracy, particularly in regions of strong gradient alignment or complex flow topology.

The present work introduces a new global convergence approach based on similar metric error energy-minimization functionals; however, minimization is done by direct displacement of the generators rather than mapping of the physical domain. By coupling global generator motion with connectivity updates, the method achieves both rapid global metric convergence and high-quality, isotropic Voronoi cells suitable for robust CFD computation.

3. Poisson-driven correction with topology regeneration

This section introduces a Poisson-based correction step that relocates the mesh generators according to the target metric field. The procedure is conceptually analogous to the pressure correction step in low-Mach-number solvers used to enforce incompressibility. The metric error to be minimized is defined as

$$e_i = \frac{A_i^{1/d} - h_i}{\min(A_i^{1/d},\, h_i)}, \tag{3.1}$$

where A_i denotes the cell area in two dimensions or the cell volume in three dimensions and d is the spatial dimension. The quantity h_i represents the target local cell size with $\mathcal{M}_i = 1/h_i^2$. The error is expressed relative to the smallest local length, ensuring that both over- and under-refined regions are penalized. This weighting is motivated by numerical

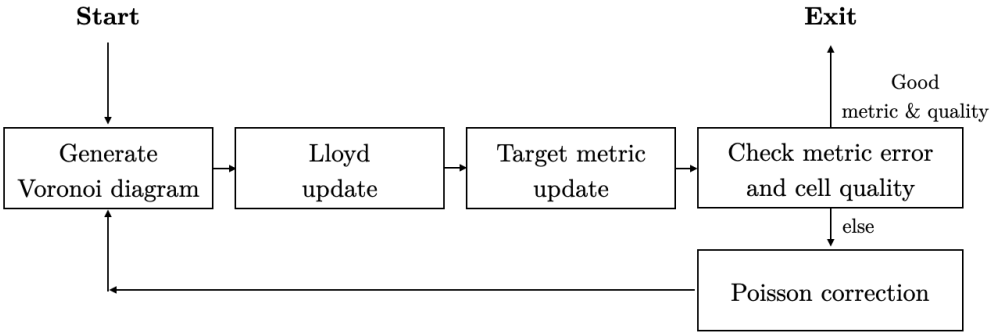

FIGURE 1: Coupled Lloyd and Poisson correction.

stability considerations: Excessively small control volumes impose restrictive time-step limitations through the Courant–Friedrichs–Lewy (CFL) condition and Fourier stability requirements.

Several formulations can be employed to minimize this error through mesh motion or deformation functionals. In the present work, the mesh adaptation is driven by a Poisson correction step. The Poisson equation is expressed as

$$\nabla^2 \phi = e, \tag{3.2}$$

where ϕ denotes a scalar potential associated with the local metric imbalance. A pseudo-velocity field is derived from the potential as

$$u_i = -\nabla \phi_i, \tag{3.3}$$

which drives the Lagrangian displacement of the mesh generators,

$$x_i^{n+1} = x_i^n + \alpha\, u_i, \tag{3.4}$$

with α a relaxation coefficient that controls the displacement amplitude. The value of α is chosen to balance convergence rate and stability, following a CFL-like restriction,

$$\alpha \max_i \|u_i\|/h_i < k, \tag{3.5}$$

where k is an empirical constant.

The mesh connectivity must be reconstructed after each Poisson correction step. However, within the Voronoi framework, this limitation is largely mitigated: Computing the current cell metric already requires regenerating the Voronoi tessellation from the updated generator positions. As a result, the cost of connectivity reconstruction is hidden within the natural regeneration step, allowing topology changes to be handled implicitly without high extra computational overhead. The exact process is shown in Figure 1. The Poisson correction step is directly coupled with the Lloyd iterations, converging at the same time quality and metric.

4. Validation on a two-dimensional test case

A two-dimensional validation case is used to demonstrate convergence toward the prescribed metric. Figure 2 illustrates the target cell size distribution. The smallest cell size is set to 0.02 within a computational domain of 2.5×1.5, with a geometric growth rate

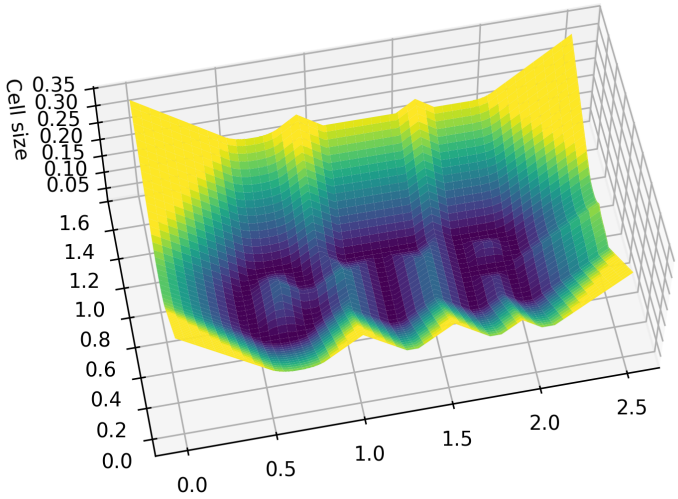

FIGURE 2: Target cell size distribution.

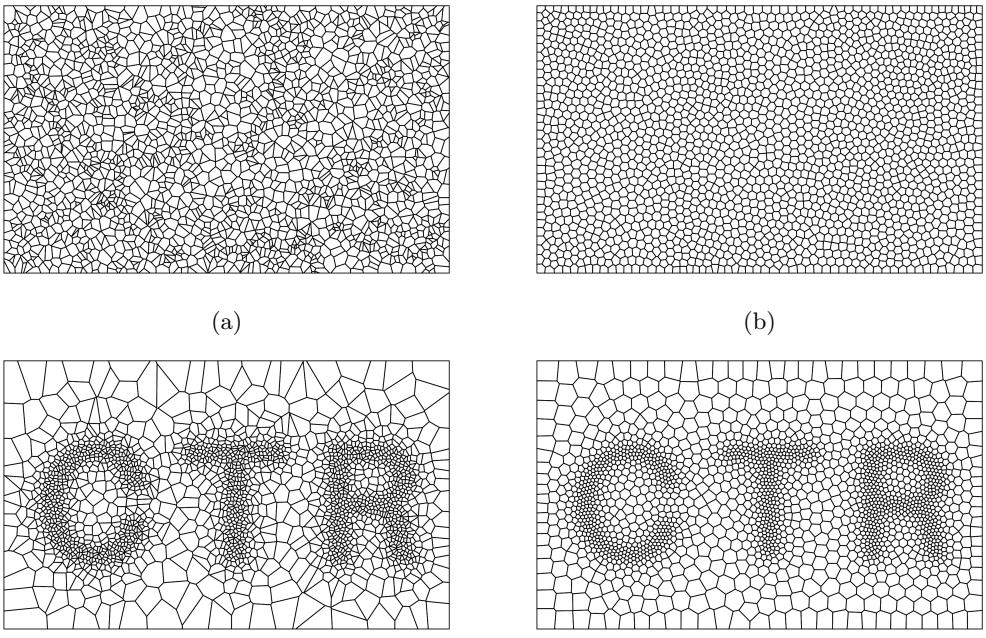

(a) (b)

(c) (d)

FIGURE 3: Resulting Voronoi diagrams; (a) Initial random uniform sampling (b) Converged Lloyd only (c) Converged Poisson correction only (d) Converged Lloyd + Poisson correction.

of 1.3. Initial generator positions are seeded randomly and their number is adjusted to match the total number of cells required by the target metric.

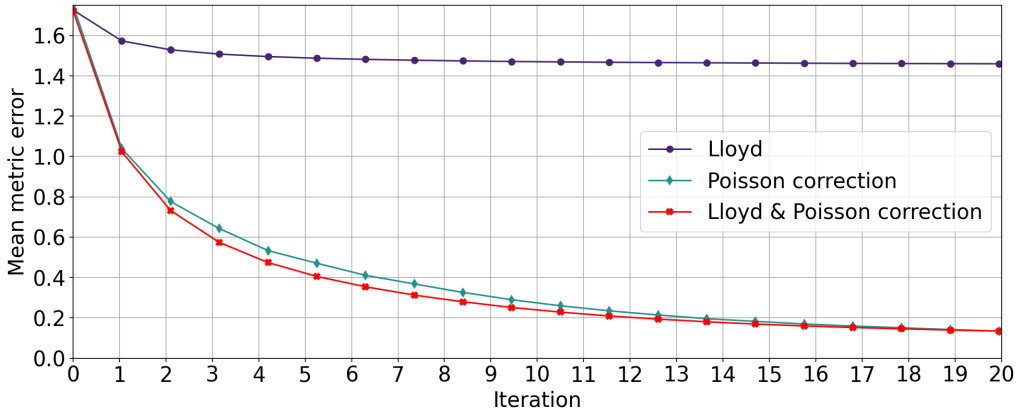

FIGURE 4: Mean metric error convergence.

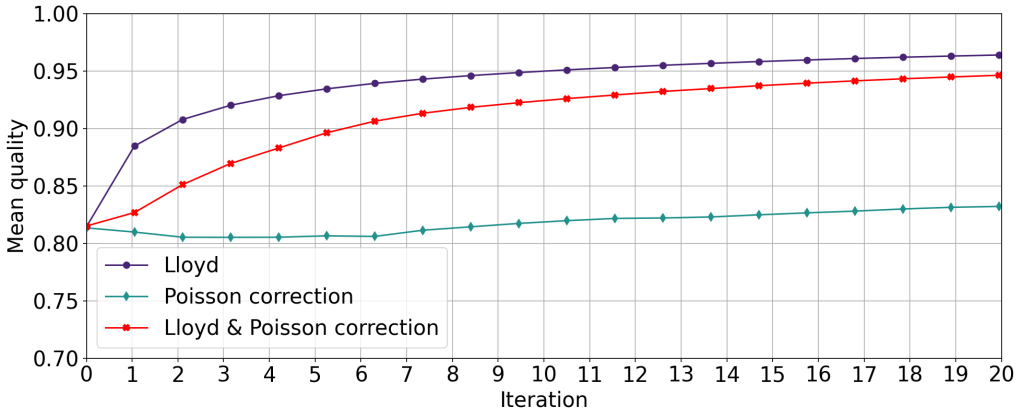

FIGURE 5: Mean quality convergence.

Figure 3 presents the obtained Voronoi diagrams produced using different mesh refinement procedures. Figure 3(a) corresponds to the initial diagram generated from random seed locations. From this initial state, three distinct adaptation strategies are applied:

- Lloyd smoothing only: The resulting Voronoi diagram exhibits excellent cell quality and a nearly uniform cell size distribution [Figure 3(b)];
- Poisson correction only: The target metric is reached, but the resulting cells remain highly irregular and poorly shaped [Figure 3(c)];
- Coupled Lloyd smoothing and Poisson correction: The resulting diagram conforms to the target metric while achieving significantly improved cell quality [Figure 3(d)].

These visual observations are quantified through two complementary metrics. The convergence toward the target metric is assessed using the error definition of Eq. (3.1); a global measure is obtained as the mean absolute error over all cells. Cell quality is evaluated from the ratio between the cell area and that of the regular n-gon having the same

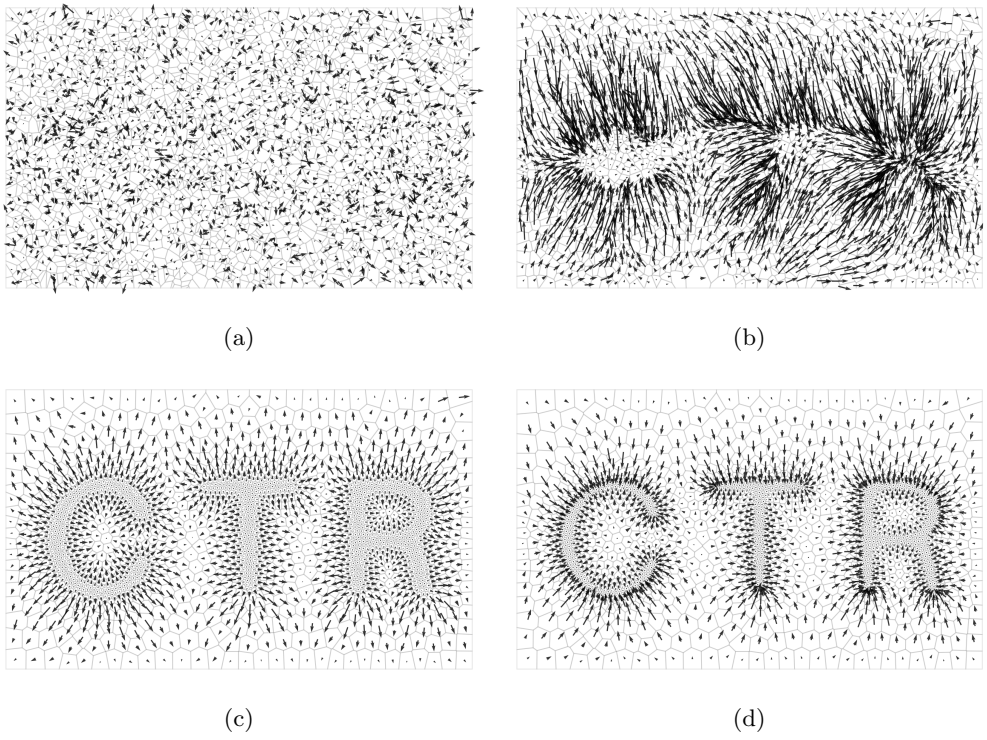

(a) (b)

(c) (d)

FIGURE 6: Generator displacements (a) Lloyd initial displacement (b) Poisson initial displacement (c) Lloyd final displacement with magnitude x5 (d) Poisson final displacement with magnitude x5.

perimeter, expressed as

$$Q = \frac{4A\,n\,\tan(\pi/n)}{P^2}, \tag{4.1}$$

where A is the cell area, P its perimeter and n the number of polygon edges. A perfectly regular cell yields $Q = 1$, while distorted cells have lower values. The evolution of both the mean metric error and the mean cell quality during the iterative process is shown in Figure 4 and Figure 5, confirming the prior observations. The immediate sharp decrease of metric error is explained by the global nature of the Poisson correction, illustrated by Figure 6(b).

The metric error decreases slightly faster when Lloyd smoothing is combined with the Poisson correction. This improvement arises because the Poisson correction is more accurate when applied at the centroid of each cell. However, the overall mesh quality slightly deteriorates when the Poisson solver is added to the Lloyd iterations. This can be explained by the presence of cell size gradients in the converged mesh, which prevent the mesh from being discretized by perfectly regular polygons. The opposition between Lloyd smoothing and the metric gradient is illustrated by Figure 6, where for the converged mesh, Lloyd and the Poisson correction perfectly balance out.

During the convergence, topology has been monitored. The results show that over 98% of the Voronoi cells undergo a topology change at some point.

5. Conclusions and perspectives

A coupled Lloyd–Poisson approach for metric-based Voronoi mesh adaptation has been proposed and evaluated. The method effectively reduces metric error while preserving satisfactory cell quality through the simultaneous action of smoothing and metric correction. Ongoing work focuses on extending this framework to large-scale three-dimensional simulations within a parallel GPU environment. This extension is particularly promising, as GPU-based Voronoi diagram generation is well established (Ray *et al.* 2018) and Poisson solvers exhibit great scalability on modern GPU architectures.

Acknowledgments

The support of ONR to CTR under grant N000142312833 is gratefully acknowledged.

REFERENCES

Agostinelli, P. W., Rochette, B., Laera, D., Dombard, J., Cuenot, B. & Gicquel, L. 2021 Static mesh adaptation for reliable large-eddy simulation of turbulent reacting flows. *Phys. Fluids* **33**, 035134.

Bénard, P., Balarac, G., Moureau, V., Dobrzynski, C., Lartigue, G. & D'Angelo, Y. 2016 Mesh adaptation for large-eddy simulations in complex geometries. *Int. J. Numer. Meth. Fluids* **81**, 719–740.

Bose, S. T., Moin, P. & Ham, F. E. 2011 Explicitly filtered large-eddy simulation on unstructured grids. *Annual Research Briefs*, Center for Turbulence Research, Stanford University, pp. 87–96.

Dapogny, C., Dobrzynski, C. & Frey, P. 2014 Three-dimensional adaptive domain remeshing, implicit domain meshing, and applications to free and moving boundary problems. *J. Comput. Phys.* **262**, 358–378.

Du, Q., Faber, V. & Gunzburger, M. 1999 Centroidal Voronoi tessellations: Applications and algorithms. *SIAM Rev.* **41**, 637–676. doi:10.1137/S0036144599352836.

Formaggia, L., Micheletti, S. & Perotto, S. 2004 Anisotropic mesh adaptation in computational fluid dynamics: Application to the advection–diffusion–reaction and the Stokes problems. *Appl. Numer. Math.* **51**, 511–533. doi:10.1016/j.apnum.2004.06.007.

George, P. L., Borouchaki, H., Alauzet, F., Laug, P., Loseille, A. & Maréchal, L. 2018 *Meshing, Geometric Modeling and Numerical Simulation 2: Metrics, Meshes and Mesh Adaptation*. Numerical Methods in Engineering Series – Geometric Modeling and Applications Set. ISTE Edition.

Habashi, W. G., Dompierre, J., Bourgault, Y., Ait-Ali-Yahia, D., Fortin, M. & Vallet, M.-G. 2000 Anisotropic mesh adaptation: Towards user-independent, mesh-independent and solver-independent CFD. Part I: General principles. *Int. J. Numer. Meth. Fluids* **32**, 725–744.

Hartmann, R., Held, J. & Leicht, T. 2011 Adjoint-based error estimation and adaptive mesh refinement for the RANS and $k-\omega$ turbulence model equations. *J. Comput. Phys.* **230**, 4268–4284. doi:10.1016/j.jcp.2010.10.031.

Huang, W., Ren, Y. & Russell, R. D. 1994 Moving mesh partial differential equations (MMPDEs) based on the equidistribution principle. *SIAM J. Numer. Anal.* **31**, 709–730. doi:10.1137/0731038.

Labelle, F. & Shewchuk, J. R. 2003 Anisotropic Voronoi diagrams and guaranteed-quality anisotropic mesh generation. In *Proc. 19th Annual Symposium on Computational Geometry (SoCG)*, pp. 191–200. ACM Press. doi:10.1145/777792.777828.

Lloyd, S. 1982 Least squares quantization in PCM. *IEEE Trans. Inf. Theory* **28**, 129–137.

Loseille, A. & Alauzet, F. 2011 Continuous Mesh Framework Part I: Well-Posed Continuous Interpolation Error. *SIAM J. Numer. Anal.* **49**, 38–60. doi:10.1137/090754078.

Mao, S. 2015 A new hybrid adaptive mesh algorithm based on Voronoi tessellations and equi-distribution principle: Algorithms and numerical experiments. *Comput. Fluids* **108**, 1–13. doi:10.1016/j.compfluid.2014.09.031.

Offermans, N., Massaro, D., Peplinski, A. & Schlatter, P. 2023 Error-driven adaptive mesh refinement for unsteady turbulent flows in spectral-element simulations. *Comput. Fluids* **251**, 105736. doi:10.1016/j.compfluid.2022.105736.

Praud, B., Breil, J. & Muscat, L. 2025 Polygonal mesh adaptation with cell size control for hypersonic flow simulations. In *Proc. SIAM Intl. Meshing Roundtable 2025*, pp. 1–13. doi:10.1137/1.9781611978575.1.

Ray, N., Sokolov, D., Lefebvre, S. & Lévy, B. 2018 Meshless Voronoi on the GPU. *ACM Trans. Graph.* **37**, 1–12. doi:10.1145/3272127.3275043.

Stock, A. & Moureau, V. 2024 Feature-based adaptive mesh refinement for multi-regime reactive flows. *Proc. Combust. Inst.* **40**, 105488. doi:10.1016/j.proci.2024.105488.