National Research Journal of Information Technology & Information Science ISSN: 2350-1278
Volume-13, Issue-1, Year-2026 (January-June) Peer Reviewed & Refereed Journal (IF: 7.9)
PP: 450-457 Journal Website www.nrjitis.in

A CONCISE REVIEW OF MESH GENERATION STRATEGIES AND QUALITY ASSESSMENT FOR NUMERICAL SIMULATIONS *PREPARED AS A REVISED REVIEW PAPER IN THERMAL AND FLUID ENGINEERING*

Kapil Sharma

Aerospace Engineer

Lalit Chhaunkar

Supervised by Shubhankar Singh, IIT Mandi M.Tech in Thermal and Fluid Engineering

Shriyansh

Supervised by Shubhankar Singh, IIT Mandi M.Tech in Thermal and Fluid Engineering

ABSTRACT

Mesh generation remains one of the most decisive stages in computational analysis because the quality of the discretized domain directly influences numerical stability, accuracy, and computational cost. In thermal and fluid engineering, the challenge is not simply to create cells that fill a geometry, but to generate a mesh that captures sharp gradients, wall-bounded behavior, curved boundaries, and localized physics without making the model prohibitively expensive. This review reorganizes the subject in a clear progression from mesh classification to quality assessment, core generation algorithms, adaptive strategies, and emerging data-driven workflows. Structured, unstructured, hybrid, and polyhedral meshes are compared in terms of flexibility, solver efficiency, and suitability for realistic engineering geometries. The discussion then links commonly used quality metrics such as skewness, orthogonality, aspect ratio, Jacobian quality, and growth rate to actual solver behavior in CFD and FEA. Classical methods including algebraic mapping, PDE-based grid generation, advancing front meshing, Delaunay refinement, and octree approaches are reviewed with emphasis on their practical strengths and limitations. The paper further examines adaptive meshing, boundary-layer treatment, immersed and overset formulations, isogeometric concepts, and the early influence of artificial intelligence on mesh decision-making. The overall argument is that robust meshing is best understood as a balance among geometry, physics, and algorithmic control rather than as a purely automated preprocessing task.

Keywords: *Mesh Generation, Computational Fluid Dynamics, Finite Element Analysis, Adaptive Meshing, Delaunay Triangulation, Boundary-Layer Meshing, Polyhedral Cells, Isogeometric Analysis*

1. INTRODUCTION

Numerical simulation converts a continuous physical domain into a finite set of cells or elements on which algebraic equations can be solved. That conversion may look like a preprocessing step, but it governs the fidelity of the entire analysis. A poor mesh can distort gradients, amplify truncation errors, slow convergence, and in severe cases make a physically correct model unusable. For this reason, experienced analysts often judge the readiness of a CFD or FEA model by the mesh long before they judge the solver setup.

The role of meshing becomes even more critical in thermal and fluid engineering because the underlying fields are rarely smooth everywhere. Boundary layers, recirculating wakes, shocks, thin passages, heat transfer interfaces, and moving boundaries all place contradictory demands on the discretization. Fine cells are needed where the solution varies rapidly, but globally uniform refinement is computationally wasteful. The central task of mesh design is therefore to distribute resolution where the physics demands it while preserving element quality and manageable run time.

This revised review presents the topic in a chronological and conceptually ordered form. It begins with mesh classes, then explains quality criteria, summarizes the principal mesh generation algorithms, discusses adaptive and application-specific strategies, and concludes with current research directions. The objective is to provide a compact but academically mature account that is useful for postgraduate work in simulation-driven engineering.

Recent literature also shows a shift in how mesh generation is evaluated. Earlier discussions often treated meshing as a separate preprocessing discipline, whereas contemporary studies increasingly judge it by its influence on solver robustness, uncertainty bounds, and engineering turnaround time. This broader view is especially relevant in research environments, where simulation credibility depends not only on elegant theory but also on repeatable computational practice.

2. EVOLUTION AND CLASSIFICATION OF MESHES

The earliest simulation grids were structured and logically rectangular. In a structured mesh, nodes are arranged in an ordered index system, so each interior node has a predictable neighbor pattern. This regularity produces compact storage, efficient matrix assembly, and solver-friendly stencils. Cartesian and curvilinear body-fitted grids belong to this family. Their weakness is geometric inflexibility: as geometry becomes more intricate, blocking and index management become labor intensive.

Unstructured meshes emerged as a response to complex industrial geometry. They store connectivity explicitly rather than implicitly, allowing triangles, quadrila terals, tetrahedra, prisms, pyramids, and polyhedra to conform to irregular boundaries. This flexibility makes them indispensable for automotive underhood flows, turbomachinery passages, biomedical domains, and detailed heat exchanger models. The trade-off is greater memory cost and somewhat lower computational efficiency per cell.

Hybrid strategies now dominate practical engineering workflows. Prism or hexahedral layers are commonly placed near walls to resolve shear and heat transfer, while tetrahedral or polyhedral cells fill the rest of the volume. Polyhedral meshes deserve special mention because they usually provide a richer neighbor stencil than tetrahedra, often improving gradient reconstruction and convergence for the same macroscopic flow problem. The modern trend is not a strict choice between structured and unstructured meshing, but rather a deliberate combination aligned with local physics and geometry.

From a historical perspective, this evolution also mirrors the growth of computing power. Early grid design favored regularity because memory and CPU budgets were severely limited. As hardware improved, mesh generators gained the freedom to support complex connectivity and localized refinement. Even so, the original structured-grid logic has not disappeared. It still informs modern multi-block methods, high-order mapping strategies, and the near-wall organization of many hybrid meshes. The field has therefore expanded by accumulation rather than replacement.

This classification is also linked to the governing numerical method. Finite difference formulations historically benefited from structured topologies, finite volume methods matured comfortably on mixed and polyhedral cells, and finite element analysis developed powerful interpolation theory on triangles, quadrilaterals, tetrahedra, and hexahedra. As a result, meshing cannot be chosen independently of the discretization framework. A mesh that is elegant for one family of solvers may be awkward or wasteful for another.

3. MESH QUALITY AND ITS NUMERICAL MEANING

Mesh quality metrics matter because they influence the numerical approximation, not merely the geometry. Aspect ratio measures elongation. In isotropic regions, excessively stretched cells tend to reduce accuracy, but in wall-bounded flow they may be entirely appropriate when the long direction follows the wall and the short direction resolves the normal gradient. Hence, aspect ratio must always be interpreted in relation to the local solution field.

Skewness is one of the most influential indicators in finite volume practice. Highly skewed cells weaken the alignment between face normals and centroid-to-centroid vectors, making flux reconstruction less reliable. Orthogonal quality describes the same issue from a different angle by quantifying how closely the cell geometry supports accurate normal-gradient evaluation. As non-orthogonality increases, the solver must rely more heavily on correction terms, which can slow or destabilize convergence.

Additional quality measures include Jacobian quality for higher-order or distorted elements, smoothness of cell-size transition, and local volume ratio. These quantities become especially important near inflation layers, rotating interfaces, and thin geometric clearances. A mesh is therefore acceptable not when every metric is globally ideal, but when each local compromise remains consistent with the physics, discretization scheme, and solver tolerances.

Quality assessment should also be tied to mesh independence. A numerically stable mesh is not automatically a sufficiently resolved one. Analysts must compare key outputs such as pressure drop, drag coefficient, wall heat flux, stress concentration, or temperature peak across progressively refined meshes. If the solution changes substantially between refineme nt levels, the model has not yet reached a defensible discretization. In that sense, quality metrics provide local geometric screening, while mesh independence provides global numerical validation.

In research reporting, this distinction is crucial. Many s imulation studies list a few quality metrics but do not explain whether those metrics were acceptable for the chosen turbulence model, interpolation order, or boundary treatment. A more convincing academic presentation connects geometry quality, solution sensitivity, and final engineering conclusions in one chain of reasoning. That is the standard expected in serious postgraduate work.

Table 1. Practical interpretation of common mesh quality metrics

Metric	Typical engineering interpretation
Aspect ratio	Prefer low values in isotropic regions; high values are acceptable in aligned boundary layers.
Skewness	Keep low for stable finite-volume convergence; large values often trigger poor flux reconstruction.
Orthogonal quality	Higher is better; low values increase correction terms and solver sensitivity.

Growth rate	Smooth transitions, often about 1.1 to 1.2 near critical regions, reduce truncation error jumps.
Jacobian quality	Must remain positive and comfortably above inversion limits, especially for higher-order elements.

4. CORE MESH GENERATION METHODOLOGIES

Classical algebraic grid generation computes interior points directly from boundary interpolation. Transfinite interpolation belongs to this family and remains useful for initialization and smoothly bounded structured blocks. Its main limitation is the lack of strong interior control. PDE-based methods improved on this by solving elliptic, hyperbolic, or parabolic systems for grid coordinates, making it possible to obtain smoother body-fitted grids with controlled clustering.

The advancing front method became influential because it builds elements outward from the boundary, preserving the quality of the surface discretization in the near-wall region. It is especially effective when boundary representation is critical. However, front collision and closure can create low-quality slivers unless the node insertion logic is carefully managed.

Delaunay triangulation and Delaunay refinement introduced a stronger geometric quality foundation. By enforcing the empty circumcircle or circumsphere property, these methods often produce more balanced elements and lend themselves well to automated tetrahedral meshing. Octree and Cartesian-dominant methods add hierarchical subdivision, which is valuable for lar ge domains, localized refinement, and high-performance computing. Many commercial platforms now combine these ideas rather than using any single methodology in isolation.

Hexahedral meshing remains a separate discussion because it promises excellent accuracy per degree of freedom but is much harder to automate for arbitrary three-dimensional geometry. Sweeping, multi-zone decomposition, and blocking-based approaches work very well when the part has favorable topology, such as ducts, blades, manifolds, or extruded passages. When the geometry loses that topological simplicity, the algorithmic difficulty rises sharply. This is one reason why tetrahedral and polyhedral workflows are still more common in general industrial preprocessing.

Another practical observation is that modern commercial software seldom exposes a single pure algorithm. A so-called tetrahedral mesher may combine octree sizing, surface recovery, Delaunay insertion, smoothing, and quality-based local retriangulation in one pipeline. Understanding the conceptual ingredients still matters, because it helps the user diagnose why the generator failed near a fillet, narrow gap, or non-manifold junction instead of treating the software as a black box.

5. ADAPTIVE, BOUNDARY-LAYER, AND EMBEDDED STRATEGIES

Static meshes assume that the analyst already knows where fine resolution is needed. Adaptive meshing relaxes that assumption by refining or relocating cells in response to the developing solution. H-refinement changes element size, p-refinement increases polynomial order, r-refinement redistributes nodes, and hp-strategies combine spatial and approximation adaptation. In transient thermal and fluid problems, these methods can reduce cost while preserving accuracy around shocks, thermal fronts, or moving interfaces.

Boundary-layer meshing remains a specialized but essential operation in viscous CFD. The first layer height, number of layers, and growth rate must be chosen from physical targets such as the desired y-plus value and expected boundary-layer thickness. A mesh that ignores this requirement may report acceptable global quality while still producing unreliable wall shear stress or heat-transfer predictions. This is why prism or hexahedral inflation layers remain standard in aerodynamic, turbomachinery, and internal-flow analysis.

Embedded approaches such as immersed boundary, cut-cell, and overset methods reduce the dependence on fully body-fitted volume meshes. They are especially useful for moving components, rapid geometry iteration, and multi-body interaction. Their attraction lies in workflow flexibility, though they introduce added complexity in interpolation, boundary enforcement, and conservation.

Adaptive meshing is also increasingly important in coupled simulations. In conjugate heat transfe r, for example, the thermal boundary layer in the fluid and the conduction path in the solid may demand very different spatial resolution. Likewise, in fluid-structure interaction or combustion, the zones of interest can move or intensify over time. A static mesh designed only from geometry is often insufficient for these cases, whereas adaptive criteria based on gradients, residuals, or a posteriori error indicators can deliver a much more rational allocation of computational effort.

The same reasoning applies in multiphase and free-surface problems. Interfaces may stretch, break, merge, or translate across the domain, and the local curvature or phase-fraction gradient can become the dominant meshing criterion. For such simulations, mesh motion and mesh ada ptation are no longer optional refinements but part of the numerical strategy itself.

6. EMERGING DIRECTIONS: ISOGEOMETRIC AND AI-ASSISTED MESHING

Isogeometric analysis was proposed to narrow the long-standing gap between CAD geometry and numerical analysis. Instead of approximating smooth CAD surfaces with piecewise linear elements, it uses spline-based descriptions such as NURBS for both geometry and field approximation. This can reduce geometric error and provide higher continuity than conventional finite elements. The main challenge is the construction of analysis-suitable volumetric parameterizations for realistic engineering models.

Artificial intelligence is entering meshing more cautiously than headlines sometimes suggest. The near-term use case is not a total replacement of geometry-based algorithms, but assistance in decisions that are difficult to tune heuristically. Learning-based size-field prediction, graph-based error estimation, and reinforcement learning for advancing-front decisions are examples of this transition. These tools are promising because they may encode expert meshing judgment in reusable form.

Even so, AI-driven meshing remains limited by training diversity, geometric robustness, and the need for physical trust. In industrial simulation, analysts still require guarantees on validity, repeatability, and conservation. The likely future is therefore hybrid: classical computational geometry for correctness, augmented by data-driven models for speed, parameter selection, and local decision support.

Another emerging direction is high-order meshing for discontinuous Galerkin and spectral-element solvers. In these methods, straight-sided linear cells can become the dominant source of geometric error even when the governing discretization is formally high order. Curved elements must therefore be generated and checked carefully to avoid negative Jacobians and unstable distortion. This requirement reconnects meshing with CAD quality, surface parameterization, and optimization-based smoothing in a way that is likely to shape future solver-preprocessor design.

A related frontier is surrogate-assisted error control. Instead of estimating the discretization error exclusively through classical residual-based procedures, researchers are exploring graph neural networks and reduced-order predictors that learn where additional resolution is likely to be useful. If these systems mature, they may shorten the refinement loop between preliminary solution and improved mesh.

Table 2. Summary of widely used meshing strategies

Method	Main strength	Main limitation
Structured / body-fitted	High solver efficiency and orderly connectivity	Difficult blocking for complex geometry
Advancing front	Strong boundary conformity	Front collision can create poor cells
Delaunay refinement	Good automation with geometric quality control	Boundary-layer handling still needs care
Octree / Cartesian-dominant	Efficient local refinement and scalability	Boundary representation may need cut-cell or hybrid treatment
Hybrid prism-tet / polyhedral	Balances wall resolution and geometric flexibility	Setup quality depends heavily on transition control

7. RELEVANCE TO THERMAL AND FLUID ENGINEERING PRACTICE

For thermal and fluid applications, mesh decisions should be tied directly to the governing transport phenomena. External aerodynamics demands careful near-wall treatment, wake refinement, and smooth far-field growth. Internal cooling passages require attention to bends, secondary flows, and conjugate heat-transfer interfaces. Boiling, combustion, or multiphase problems may require localized refinement near phase boundaries, reaction zones, or interfacial transport regions.

A practical workflow begins with geometry cleanup and a preliminary topology choice, followed by targeted local sizing, boundary-layer planning, and an initial quality audit. The first simulation run should not be treated as final validation but as a diagnostic stage that reveals where the mesh under-resolves the solution. Mesh independence studies, residual behavior, monitored integral quantities, and physically meaningful field checks must all inform subsequent refinement.

From an engineering management perspective, the best mesh is rarely the smallest or the most elegant. It is the mesh that delivers dependable predictions within the available computational budget and project timeline. That practical viewpoint is especially important for postgraduate work, where methodological rigor must be balanced against reproducibility and resource constraints.

Common commercial and open-source tools reflect this philosophy in different ways. ANSYS Meshing emphasizes workflow automation with user-controlled local sizing, inflation, and quality diagnostics. ICEM CFD remains valuable where blocking discipline and structured control are required. Open FOAM combines block Mesh for background topology and snappy Hex Mesh for feature-driven refinement around CAD surfaces. STAR-CCM+ integrates polyhedral, trimmed, and overset capabilities into a comparatively unified environment. Knowing t he logic of these platforms helps the analyst choose the tool chain that best matches the physics and geometry of the problem.

For students and early researchers, this software awareness is important because different tools encourage different meshing habits. A structured-grid background often teaches discipline in topology and boundary alignment, whereas automated unstructured tools encourage rapid iteration and local refinement. The strongest analysts usually learn from both traditions and combine them selectively.

Validation should finally connect the mesh back to the engineering question. In a thermal-fluid study, that may mean demonstrating that velocity, pressure, temperature, and wall quantities have all reached acceptable sensitivity limits with respect to mesh refinement. In other words, a mesh is not justified because it looks neat in a preprocessor window; it is justified because the reported physical conclusions remain stable when the discretization is challenged.

8. CHALLENGES AND FUTURE SCOPE

Several important problems remain open. Fully automatic, high-quality all-hexahedral meshing for arbitrary three-dimensional geometry is still not a solved industrial capability. High-order curved mesh generation also remains difficult because geometric fidelity must be preserved without creating negative Jacobians or unstable distortion. In large eddy simulation and direct numerical simulation, the absolute scale of the mesh itself becomes a computational bottleneck, pushing interest toward GPU-aware generation pipelines and multiscale adaptivity.

Another unresolved issue is interoperability between CAD systems, mesh generators, and solvers. Real engineering geometry often contains gaps, slivers, tiny features, and non-manifold regions that are tolerable in design software but troublesome in simulation. Progress in robust geometry repair, automatic feature recognition, and analysis-suitable parameterization is therefore as important as progress in the meshing algorithms themselves.

The future of mesh generation will likely be defined by integration rather than replacement. Better geometry handling, adaptive automation, high-order discretization support, and data-assisted meshing are all converging toward the same goal: making simulation-ready meshes faster to obtain without reducing physical credibility.

For researchers, another challenge lies in evaluation methodology. New meshing algorithms are sometimes compared only by cell count or a small set of geometric quality metrics, whereas the more meaningful benchmark is the quality of the final numerical prediction per unit computational cost. Future work should therefore connect meshing research more tightly with solver convergence history, uncertainty quantification, and application-specific validation data.

This broader perspective also encourages ethical use of automation. As meshing becomes more push-button, there is a risk that users may accept solver-ready grids without understanding the compromises built into them. Advanced tools should therefore reduce routine labor while still making assumptions visible to the user.

CONCLUSION

Mesh generation is not a routine technical formality but a central modeling decision that shapes the quality of every numerical result. The evolution from structured grids to hybrid, polyhedral, embedded, and data-assisted strategies shows that no single mesh paradigm is universally superior. Each approach offers a different balance of geometric flexibility, solver efficiency, and control over numerical error.

For thermal and fluid engineering, the most defensible meshing practice is one that connects discretization choices to actual transport physics, validates them through quality metrics and mesh-independence checks, and remains conscious of computational cost. When that balance is maintained, mesh generation becomes not a source of uncertainty, but a disciplined foundation for trustworthy simulation.

Accordingly, a strong review of meshing should not treat the subject as a catalog of software buttons or isolated algorithms. It should treat mesh generation as a synthesis of geometry processing, numerical analysis, and engineering judgment. That synthesis is what ultimately allows simulation to become predictive rather than merely illustrative.

REFERENCES

1. Bern, M., and Plassmann, P. (1999). Mesh Generation. In Handbook of Computational Geometry. Elsevier.
2. Berger, M. J., and Oliger, J. (1984). Adaptive Mesh Refinement for Hyperbolic Partial Differential Equations. Journal of Computational Physics, 53(3), 484-512.
3. Blacker, T. D., and Stephenson, M. B. (1991). Paving: A New Approach to Automated Quadrilateral Mesh Generation. International Journal for Numerical Methods in Engineering, 32(4), 811-847.
4. George, P. L., and Seveno, E. (1994). The Advancing Front Mesh Generation Method Revisited. International Journal for Numerical Methods in Engineering, 37(21), 3605-3619.
5. Hughes, T. J. R., Cottrell, J. A., and Bazilevs, Y. (2005). Isogeometric Analysis: CAD, Finite Elements, NURBS, Exact Geometry and Mesh Refinement. Computer Methods in Applied Mechanics and Engineering, 194, 4135-4195.
6. Lo, S. H. (2002). Finite Element Mesh Generation and Adaptive Meshing. Progress in Structural Engineering and Materials, 4(4), 381-399.
7. Mittal, R., and Iaccarino, G. (2005). Immersed Boundary Methods. Annual Review of Fluid Mechanics, 37, 239-261.
8. Peskin, C. S. (1972). Flow Patterns around Heart Valves: A Numerical Method. Journal of Computational Physics, 10(2), 252-271.
9. Ruppert, J. (1995). A Delaunay Refinement Algorithm for Quality Two-Dimensional Mesh Generation. Journal of Algorithms, 18(3), 548-585.
10. Shewchuk, J. R. (1998). Tetrahedral Mesh Generation by Delaunay Refinement. Proceedings of the Annual Symposium on Computational Geometry, 86-95.
11. Si, H. (2015). TetGen, a Delaunay- Based Quality Tetrahedral Mesh Generator. ACM Transactions on Mathematical Software, 41(2), Article 11.
12. Thompson, J. F., Thames, F. C., and Mastin, C. W. (1974). Automatic Numerical Generation of Body - Fitted Curvilinear Coordinate System for Two-Dimensional Bodies. Journal of Computational Physics, 15(3), 299-319.