Part IV
Finite Elements in 3-D

To define useful basis functions, one must first have a proper mesh. Consider a three-dimensional domain, convex or nonconvex. To approximate it well, design a mesh of disjoint (nonoverlapping) tetrahedra. In numerical analysis, these are called finite elements [9, 69].

In the mesh, we have nodes: the corners of the tetrahedra. The node is the most elementary ingredient in the mesh. An individual node may serve as a corner in a few tetrahedra.

If they belong to the same tetrahedron, then the nodes must be connected to each other by an edge. An edge is often shared by a few adjacent tetrahedra. There are two kinds of edges: a boundary edge could be shared by two tetrahedra. An inner edge, on the other hand, must be shared by more tetrahedra.

Each tetrahedron is bounded by four triangles: its sides or faces. In the mesh, there are two kinds of faces: an inner face is shared by two adjacent tetrahedra. A boundary face, on the other hand, belongs to one tetrahedron only.

How to construct the mesh? Start from a coarse mesh that approximates the domain poorly, and improve it step by step. For this purpose, refine: split coarse tetrahedra. At the same time, introduce new (small) tetrahedra next to the convex parts of the boundary, to improve the approximation from the inside. This procedure may then repeat time and again iteratively, producing finer and finer meshes at higher and higher levels.

This makes a multilevel hierarchy of finer and finer meshes, approximating the original domain better and better. In the end, at the top level, those tetrahedra that exceed the domain may drop from the final mesh. This completes the automatic algorithm to approximate the original domain well.

The mesh should be as regular as possible: the tetrahedra should be thick and nondegenerate. Furthermore, the mesh should be as convex as possible. Only at the top level may the mesh become concave again. A few tricks are introduced to have these properties.

To verify accuracy, numerical integration is then carried out on the fine mesh. For this purpose, we use a simple example, for which the analytic integral is well-known in advance. The numerical integral is then subtracted from the analytic integral. This is the error: it turns out to be very small in magnitude. Furthermore, our regularity estimates show that the mesh is rather regular, as required.

Once the basis functions are well-defined in the fine mesh, they can be used to approximate a given function, defined in the original domain. This is indeed the spline problem [5, 11, 17, 28, 41, 51, 58, 75, 76]: design a smooth piecewise-polynomial function to “tie” (or match) the original values of the function at the mesh nodes.

The spline problem could also be formulated as follows: consider a discrete grid function, defined at the mesh nodes only. Extend it into a complete spline: a smooth piecewise-polynomial function, defined not only at the nodes but also in between. The solution must be optimal in terms of minimum “energy.” This is indeed the smoothest solution possible.

Our (regular) finite-element mesh could be used not only in the spline problem but also in many other practical problem. Later on, we’ll see interesting applications in modern physics and chemistry.