Fitting low-dimensional manifolds to high-dimensional point clouds is among the most fundamental practices in contemporary data analysis. The underlying assumption behind the success of all such techniques is that the points are well-approximated by a single low-dimensional manifold. This talk will describe a topological technique for automatically detecting data points which fail to satisfy this manifold hypothesis. As a consequence, it becomes possible to decompose the data points into different clusters, each of which can be well-approximated by a single manifold. This is joint work with B Stolz, J Tanner and H Harrington.