Locating Buildings in Aerial Photos

Jim Green
Department of Computer Science
Moravian College
Bethlehem, PA 18018

Algorithms and techniques for use in the identification and location of large buildings in digitized copies of aerial photographs are developed and tested. The building data would be used in the simulation of objects located in the vicinity of an airport that may be detected by aircraft radar. Two distinct approaches are considered.

Most building footprints are rectangular in form. The first approach studied is to search for right-angled corners that characterize rectangular objects and then to connect these corners to complete the building.

This problem is difficult because many non-building objects, such as street corners, parking lots, and ballparks often have well defined corners which are often difficult to distinguish from rooftops. Furthermore, rooftops come in a number of shapes, sizes, shadings and textures which also limit the discrimination task.

The strategy used linear sequences of difference samples to detect straight edge segments at multiple angles and to determine when these segments meet at approximately right-angles with respect to each other.

This technique is effective in locating corners. The test image used has a fairly rectangular block pattern oriented about thirty degrees clockwise from a vertical and alignment, and the overall measurement data reflect this. However, this technique does not discriminate between buildings and other objects at an operationally suitable rate. In addition, since multiple paths are tested for each image pixel, this is a time consuming task. The process can be speeded up by preprocessing the image to locate the more optimal sampling paths.

The second approach is to rely on an human operator to identify and select the building objects and then to have the computer determine the outline and location of the selected structures.

When presented with a copy of a digitized aerial photograph, the operator uses a mouse and cursor to select a target building.

After a button on the mouse is pressed, with the cursor fully within the perimeter of the building, the program scans from the position of the cursor to a perimeter position where a shift in grayscale is detected. Once at the perimeter, the process traces along it until it, around the building, until it eventually returns to the perimeter starting point.

Spatial resolution limits cause the perimeter trace to be somewhat course so that a line straightening algorithm is employed. One result is that the building corner positions become more distinctly defined.