International Cartographic Association
Association Cartographique Internationale

Weighted Median Filtering for Terrain and Contour Generalization

Leland Brown
Independent Researcher
leland@textureshading.com

Abstract: When working with terrain elevation data, as in image processing, it is often desirable to smooth the data, for purposes such as map generalization or removal of noise. Median filtering is one common technique that can be used for this purpose. It differs from linear filtering techniques like local averaging or Gaussian blurring by its ability to smooth while retaining sharp edges in an image. When applied to elevation data, this means that median filtering can better preserve steep slopes and cliffs while otherwise reducing noise or excessive detail in the terrain.

However, median filtering as typically applied can also introduce new artifacts, such as lopping off the tops of peaks and ridges to create flat plateaus that don't exist in the original landscape. A lesser known technique, a weighted median filter, can reduce or eliminate these artifacts. This method shows promise as a way to generalize digital elevation models, as well as their associated contour lines. It can also be used to smooth hillshaded images, preserving the sharp transition in shading that occurs across ridges. And due to its ability to retain discontinuities in the data, it can be used to locate latent terracing effects hidden in elevation data, which may represent real terrain features or may indicate artifacts of the processing methods used to generate the data.

Keywords: weighted median, relief mapping, digital elevation models

1. Median vs. Weighted Median

A standard median can be conceptualized as in Figure 1(a). After sorting a set of values into numerical order, the median is then the value in the middle of the list (if there are an odd number of values). Algorithms exist to find the median without actually performing the sorting, but the result is the same.

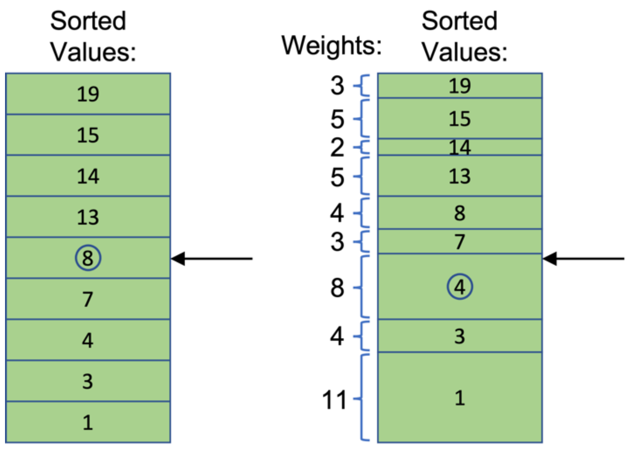

(a) (b)

Figure 1. Median (a) vs. weighted median (b).

In contrast, a weighted median can be described by imagining the sorted values as a stack of boxes, with the height of each box representing a weight for that value—as in Figure 1(b). Then the weighted median is the value in the box located at the midpoint of the height of the stack. Values with larger weights have more influence on which value is chosen as the weighted median. In this example, the set of values is the same in cases (a) and (b), but since the larger weights are mostly on the smaller values, the weighted median is "pulled" down to a lower value closer to where the bulk of the weights are (in this case, from 8 down to 4).

(We will not consider here the case that the midpoint falls exactly on the boundary of two boxes, since we can always avoid this situation by choosing integer weights whose sum is odd. For a 2D filter we can accomplish this by simply choosing weights with 4-way symmetry and an odd weight for the center pixel.)

The idea of weighting the median filter is in some ways analogous to using a weighted averaging filter, though the two types of filters have very different properties. In particular, as we will see, a weighted median retains the filter's ability to preserve sharp transitions in the data. This is because a median filter always selects one of the values existing in the original data (whereas an averaging filter combines them to compute a new value). This property does not change with a weighted median.

When using the median as a filter on image data or gridded terrain data, the procedure is as follows: for each pixel, consider a window (typically square or circular) centered on that pixel, then compute the median of the pixel values in that window, and use the resulting median as the value for that pixel in the new image or data grid.

Proceedings of the International Cartographic Association, 6, 4, 2024. https://doi.org/10.5194/ica-proc-6-4-2024
12th Mountain Cartography Workshop of the ICA Commission on Mountain Cartography, 11–15 April 2023, Snow Mountain Ranch, CO, USA.
This contribution underwent single-blind peer review based on submitted abstracts.

With the weighted median we do the same, but with higher weights on pixels near the center of the window (i.e., near the pixel under consideration) and weights dropping off to zero as we move away from the center. In this way, as we slide the window pixel by pixel across the image, pixels at the edges entering and exiting the window are less likely to affect the result, until they come closer to the window center. This, in turn, generates smoother transitions in certain cases, while retaining many features of a standard median filter.

The output of this technique is driven by the choice of weighting scheme—including the size of the kernel window, its footprint (square, circular, etc.), its symmetries, and the exact shape of the distribution of weights. In this paper, I will only consider weighting schemes that are radially symmetric, since there is generally no need to introduce an anisotropic component to the filtering process. Some trial and error has shown that a two-dimensional Gaussian function offers one reasonable choice for selecting the weights. (However, with weighted median filtering there is nothing unique about this particular function—its utility just comes from its smooth shape and symmetry, along with being highest in the center and dropping off gradually to zero.) This leaves the width of the Gaussian kernel as the remaining parameter that we can use to adjust the degree of smoothing.

2. Effects on Terrain Features

We can compare the effects of different filters by considering three filtering techniques: Gaussian smoothing (a kind of weighted average), a median filter, and a weighted median filter (using Gaussian weights). We can apply these to two prototypical terrain features: a cliff and a ridge (or peak). Figure 2 shows a series of profile views, where I have taken a vertical slice through the prototypical terrain. In each case the original terrain shape is shown, followed by the result of the filtering.

In Figure 2(a) and (b), we can see that the Gaussian weighted average smooths the terrain nicely, but it softens the cliff into a sloped hill and also eliminates the sharp top of the ridge, lowering its height. On the other hand, the median filter in (c) preserves the cliff perfectly. In (d), however, it not only lowers the top of the ridge but introduces new corners, as if the top of the ridge were sliced off. These corners, or slope discontinuities, will become visible artifacts if this terrain is rendered as a hillshaded relief image.

The weighted median, in Figure 2(e), also preserves the cliff. And in (f), although it reduces the top of the ridge like the other filters, it returns to a smoother shape than the standard median produced.

A close examination of Figure 2(f) shows a jagged stair-step effect at a fine scale. This is because the weighted median always selects from among the existing set of discrete elevation values in the original data—there is no averaging or interpolating. This occurs independently of the weights used, whether integers or real numbers. Thus, it may be desirable to apply a small amount of additional smoothing using a weighted average technique. Another possible solution would be to develop a “fuzzy” weighted median that interpolates in some fashion between the two values closest to the middle of the stack of weights.

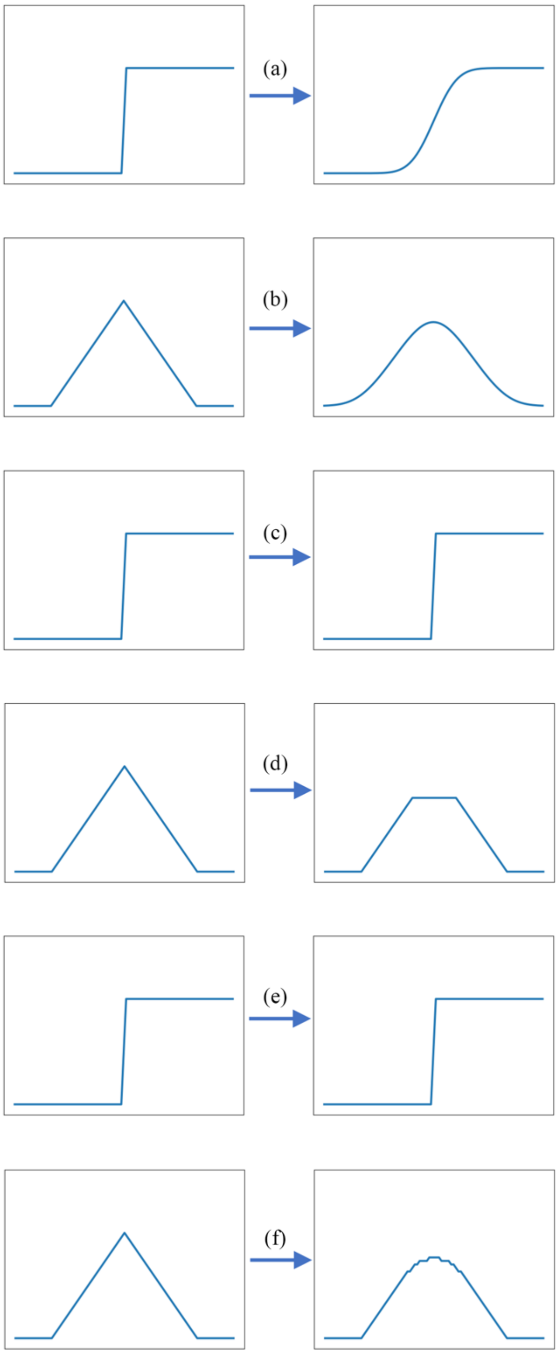

Figure 2. Profile view effects of Gaussian smoothing (a) and (b), median filter (c) and (d), and weighted median (e) and (f) using Gaussian weights.

Next, we turn to a plan view in order to get a more three-dimensional perspective and to see how the various filters affect contour lines of the terrain. Consider a cliff with a right-angle corner as in Figure 3(a), forming a flat plateau

Proceedings of the International Cartographic Association, 6, 4, 2024. https://doi.org/10.5194/ica-proc-6-4-2024
12th Mountain Cartography Workshop of the ICA Commission on Mountain Cartography, 11–15 April 2023, Snow Mountain Ranch, CO, USA.
This contribution underwent single-blind peer review based on submitted abstracts.

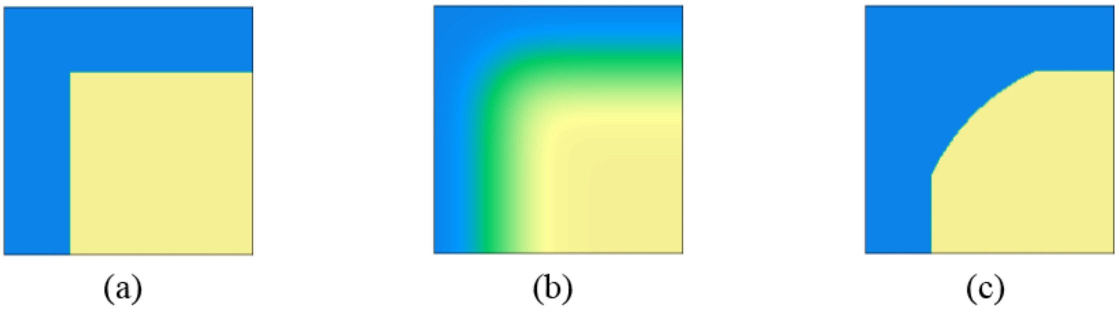

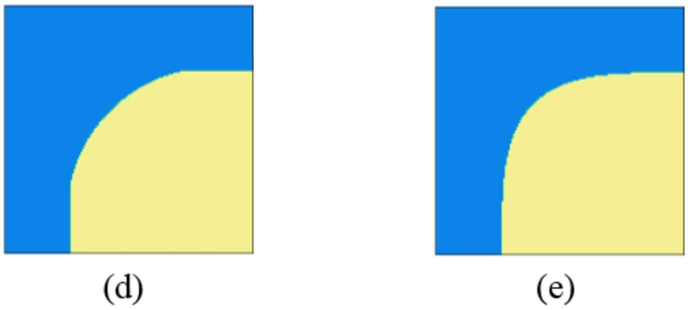

Figure 3. Plan view effects: original plateau with corner (a), Gaussian smoothing (b), median filter with square kernel (c) and circular kernel (d), and weighted median (e).

above a flat valley—or perhaps forming a section of coastline. The colors represent hypsometric tints showing different elevations.

As expected, the Gaussian smoothing in Figure 3(b) smooths out the cliff and separates the contours at intermediate heights; these are generally undesirable effects when smoothing terrain, since a cliff is an important feature.

Figures 3(c) and (d) show results of a median filter with different window shapes – square and circular, respectively. Both preserve the cliff while smoothing the sharp corner in the elevation contours, but the square filter cuts the corner rather abruptly. The circular median is more acceptable; and the weighted median is similar, or perhaps a bit more cleanly rounded, as seen in (e).

3. Examples with Real-World Data

The examples below are all taken from a region in the San Gabriel Mountains in California, USA. The data is based on the National Elevation Dataset.

Figure 4(a) shows the original digital elevation model (DEM) used here, rendered using a standard hillshading technique.

Figure 4(b) was generated by first applying a Gaussian smoothing filter to the DEM and then rendering the new DEM using hillshading. All the abrupt transitions in the terrain are smoothed out, and the image almost appears blurry. The ridges appear as bright bands because they have been flattened somewhat.

Figure 4(c) is similar, but using a circular median filter on the DEM instead. The ridge tops and canyon bottoms appear more flattened, as discussed earlier with Figure 2(d).

Figure 4(d) replaces the median filter with a weighted median filter; the ridges and canyons are smoother. Terracing artifacts are apparent as alternating light and dark bands, and they become more obvious with larger filter kernels. These may be processing artifacts from the original data, which are then preserved by the filter as other details are filtered out.

Figure 4(e) is different: in this case I've taken the original image (a) and applied the weighted median to the grayscale image pixels (not the DEM data), *after* rendering the hillshade.

In Figure 4(f) I have applied a Gaussian smoothing to the DEM, rendered an image, and then applied a weighted median to the image. The weighted median reduces the brightness of the ridges and further generalizes the image. Since two filters are being applied here, I use smaller windows for each than with the earlier examples.

Source code for software to perform a weighted median filter can be found at https://app.box.com/v/weighted-median-filter/.

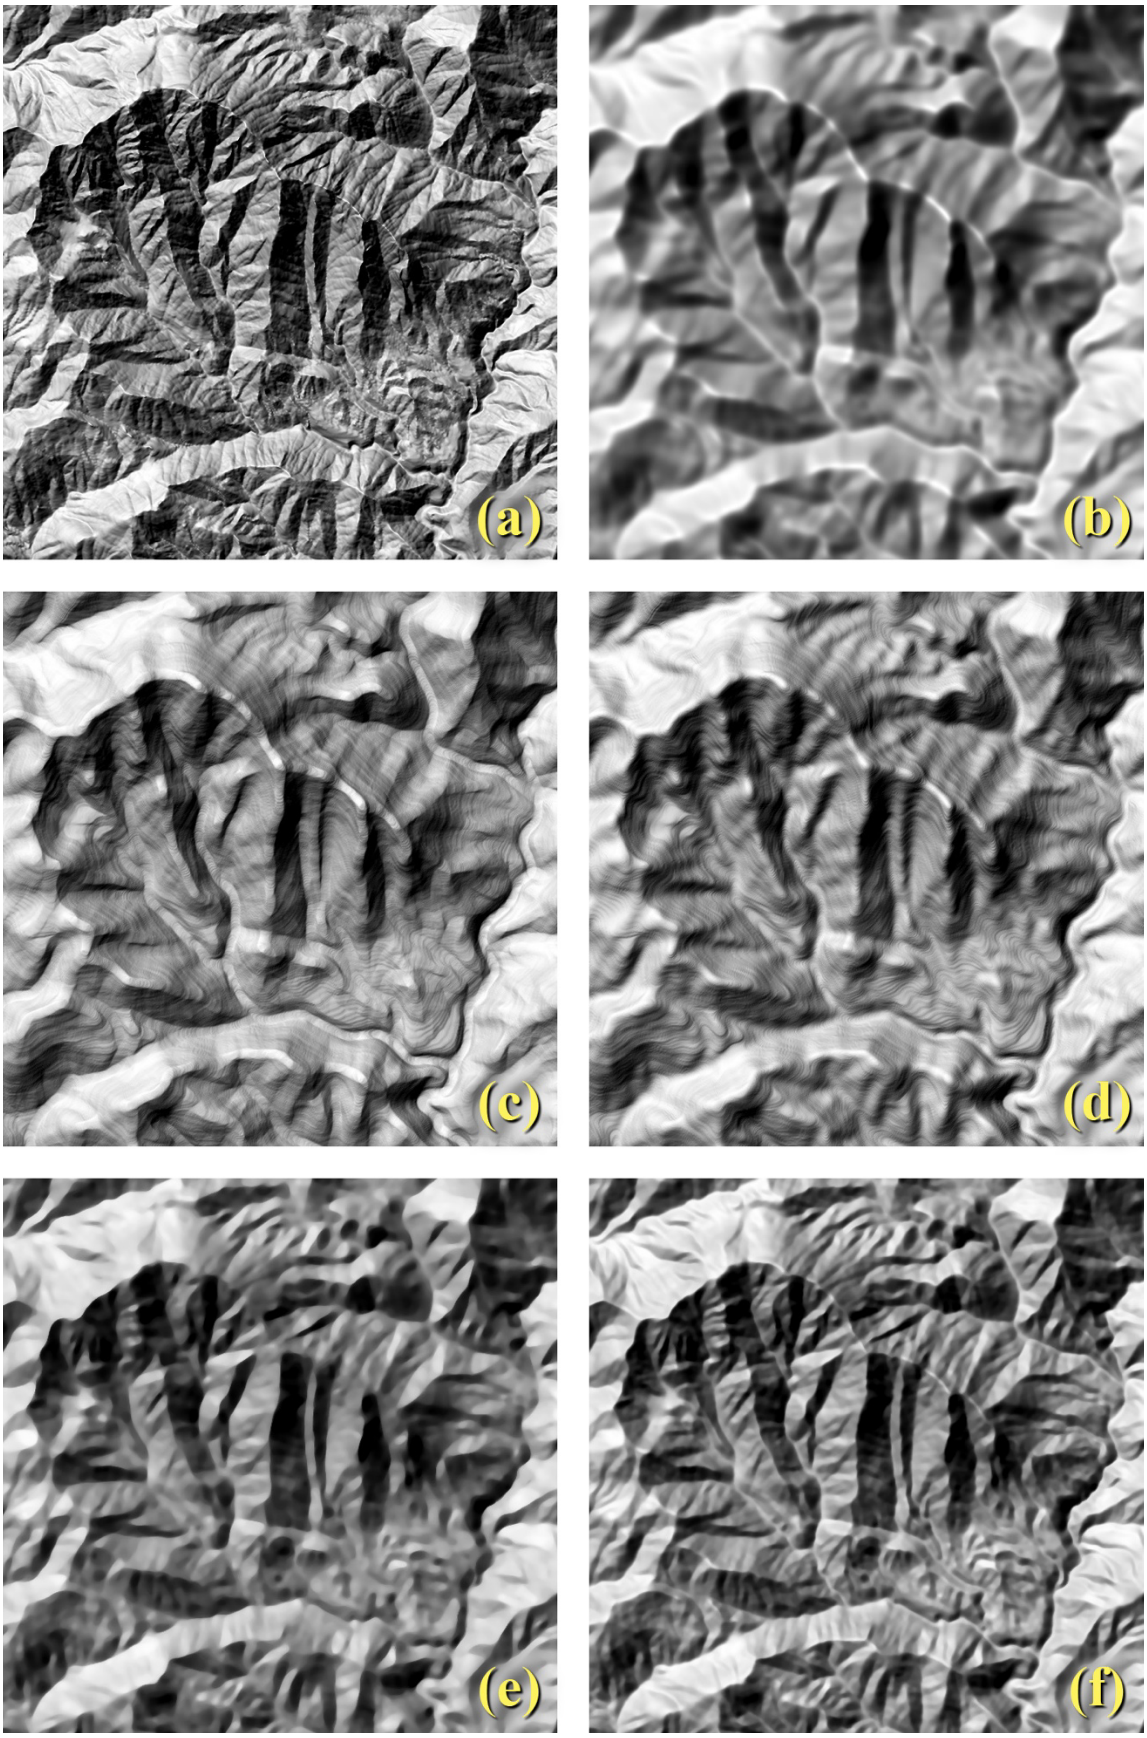

Figure 4. Original DEM (a), with Gaussian smoothing (b), median filter (c), weighted median (d), weighted median applied to original image (e), Gaussian smoothing of DEM with weighted median applied to resulting image (f).

Proceedings of the International Cartographic Association, 6, 4, 2024. https://doi.org/10.5194/ica-proc-6-4-2024
12th Mountain Cartography Workshop of the ICA Commission on Mountain Cartography, 11–15 April 2023, Snow Mountain Ranch, CO, USA.
This contribution underwent single-blind peer review based on submitted abstracts.