Noise Reduction for the Psychopath: Controlling the Blur

Chapter 10 is very long even without this fragment, which appeared in the first draft. If you're seriously interested in the theory of noise reduction you may wish to give it a look. The beta readers felt, and I agree, that the number of people who might find it useful is dwarfed by that of people who would be baffled and frustrated by it. So, out it came.

—D.M.

Postcard 10.15 *(opposite) A typical day scene, according to my experience, in the inaptly named Sunshine Coast of British Columbia, Canada.*

setting up a custom overlay rather than just using a mindless action, it surely is worth it to see what happens with different blurs, particularly when amorphous shapes like these clouds are present.

When Blurring Can Sharpen

Cloud formations are more sensitive to changes in the blurring settings than are, say, waterfalls. Assume, please, that Gaussian Blur is the only filter available. It gets applied to the inverted blue overlay. Figures 10.16B, 10.16C, and 10.16D show the impact of Radii of 25, 40, and 55 pixels, respectively. (To put these numbers in context, the printable width in Postcard 10.15 is slightly less than 2000 pixels; larger files require larger Radii to achieve the same effect.)

The decision of which of the three is best is hereby left to you. It doesn't take long to preview different blurs, once you've decided to experiment, even if you aren't using the panel.

There is no denying that blurring the overlay is a necessity. You may rightly, however, wonder why this is so, and whether there are negative consequences.

A half-real, half-artificial example should suffice. Figure 10.17A shows an upsized piece of a section of Niagara Falls. This kind of picture is tailor-made for the Bigger Ham-

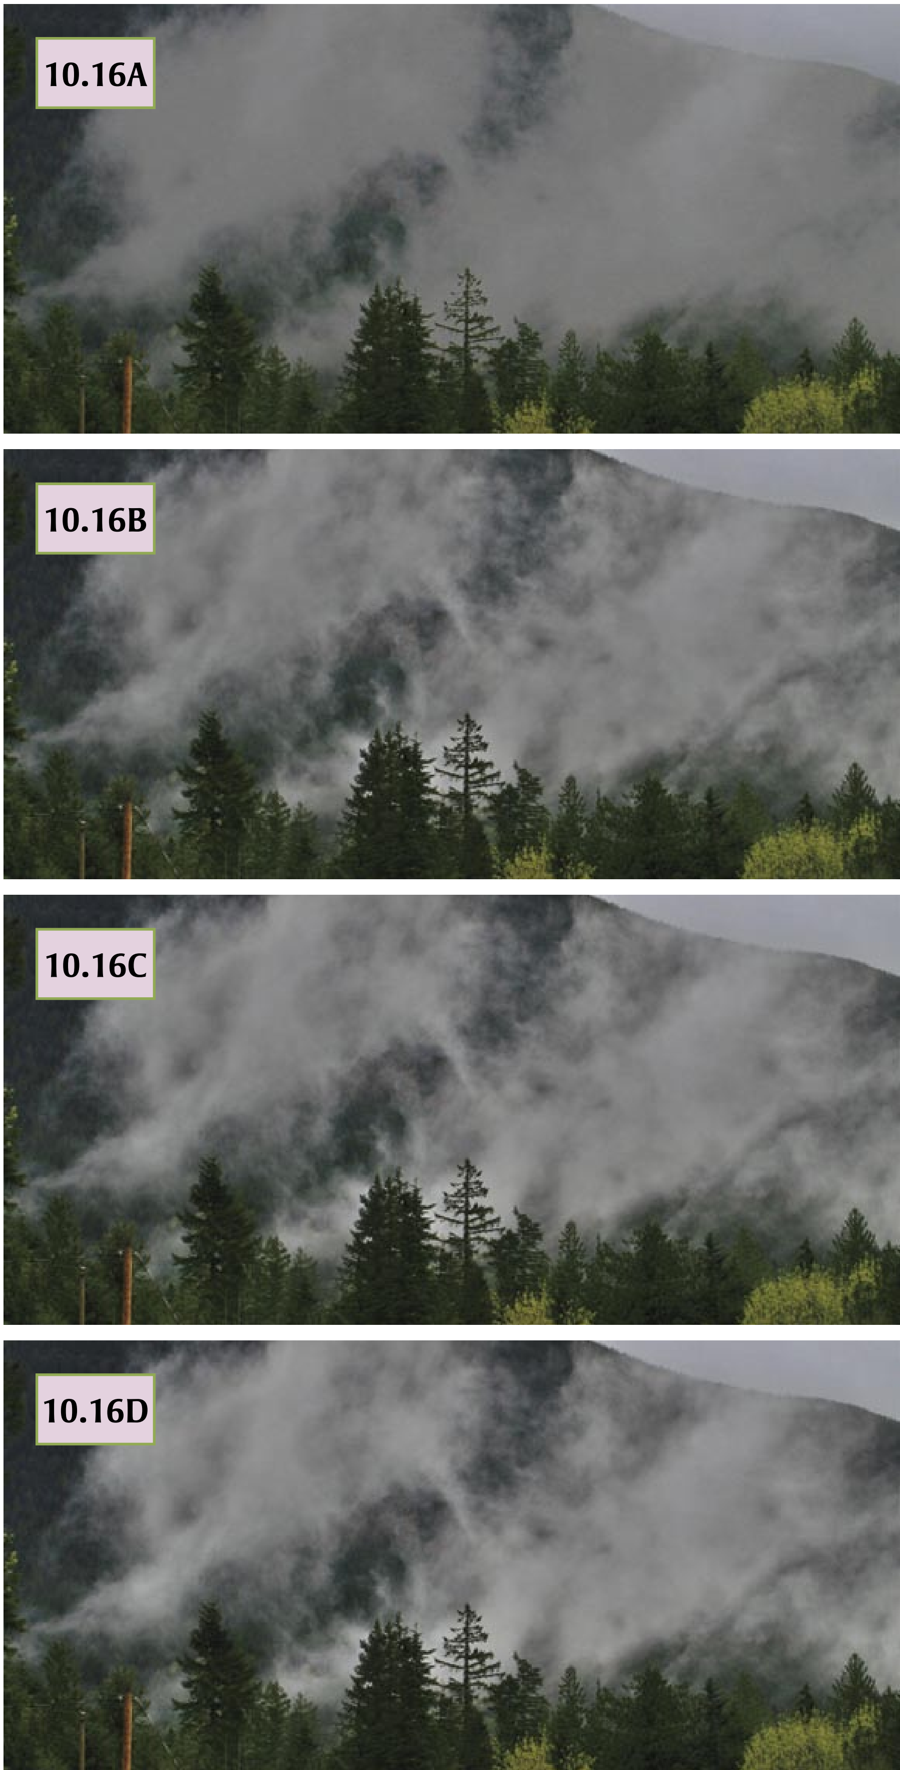

Figure 10.16 *Top, an inverted copy of the blue channel is overlaid onto the original. The remaining three versions blur it. Radius increases each time.*

mer. The top half contains natural detail. In the bottom half I have added a generous amount of dark noise, which I ask you to take as detail of a different species. The real desirable detail in this picture is the shadows in the falling water. These shadows are much wider than the noise.

Figure 10.17B has an inverted copy of itself on an overlay layer. There's no darkening layer, no blurring. Since the whole picture is quite neutral, using a single inverted channel as the overlay would probably be indistinguishable.

In the cloudy day we just worked on, at this point we were at Figure 10.16A, which was worse than the original. The same cannot be said of the noise-free areas of Figure 10.17B. They show more detail in the water than the original, Figure 10.17A, does. Without the blur, though, it's not enough.

Those with insatiable curiosity about such trivia can find the formula for computing an overlay in the Notes & Credits section. Those who do not wish their eyes to glaze over can treat overlay as a kind of S-shaped curve. The extreme highlights and shadows are short-changed, while everything between moves more toward a midtone.

Therefore what we perceive as extra contrast in the water of Figure 10.17B isn't quite what we might think. The whitest water gains little. But the darker water gets significantly darker, and we take that for extra detail.

Meanwhile, the difficulty seen in the clouds

Figure 10.17 *Top left, an area of artificial noise is added to an original photo of Niagara Falls. Top right, an inverted copy of the RGB is overlaid onto the original. Bottom left, the overlay gets a narrow-radius blur, bringing out more of the noise but not doing much for the background. Bottom right, a wide-radius blur is substituted.*

of Figure 10.16A has been avoided. When an image has significant detail both above and below the midtone, the overlay closes it up. That was the case with the aforementioned clouds, but not with this water, which is entirely lighter than a midtone. We notice midtone detail loss in the right side of the noisy area.

Figure 10.17B's overlay, being inverted, is dark in both noise and whitewater—but the water is the darker of the two. Therefore, when the overlay does its work, the water gets darker faster than the noise does; the contrast between them lessens. For the same reason, the shadows in the water do not deliver the punch they should.

The solution is to obliterate the differences in the overlay layer. Interestingly, since the noise is much smaller than the true detail, the obliteration, in the form of a blur, can be selective.

A relatively small blur, 5 pixels for example, wipes out the noise in the overlay and therefore accentuates it in Figure 10.17C. It does not, however, accomplish much with the darker water. For that, we'd need something like the 50-pixel blur that produced Figure 10.17D. In it, the noise is not as bad as in Figure 10.17C, but the water is better.

The Role of the Blur

Professional color correction favors the aggressive. Nobody I know has gotten rich in the field by overestimating the taste of typical clients. If you don't know your audience's taste, you should go for bright colors and heavy sharpening. I've had this debate elsewhere and don't intend to pursue it here, but one point about our current techniques is somewhat new.

The rules for viewer acceptance of our work are as simple as they are stupid and meaningless. Bright colors are always right, except when they are too much. Added contrast is highly preferred, as long as it doesn't go too far. Strong sharpening is also a great idea, as long as it isn't too heavy. Now all we have to do is guess what *too much* and/or *too far* and/or *too heavy* mean to people we don't know.

Your judgment as to whether the color is too vivid is apt to be warped if you insist on evaluating it next to some of the impossibly loud products of the Color Boost action. Avoid that error, though, and you're looking at the same thing another viewer would. I don't believe that your awareness that the MMM action was used to produce the color would skew your opinion.

The same is even more true of luminosity blends and curves. You perceive, more or less, what the viewer does.

Starting with this chapter and continuing through four others, however, a nasty complication introduces itself. We are using large blurs as masks to alter the perception of contrast and depth. It is related to what is known generically as unsharp masking, but traditional USM uses much smaller blurs, because bigger ones were impossible technically at the time the methods were developed. Today there is no such limitation; we use a nontraditional form of USM known as hiraloam. We also use something closely related in the blurred multiplication masks shown in the next three chapters. The Bigger Hammer is a more distant relative, and the Shadows/Highlights command even farther removed.

This is the first of five chapters that are thus seriously concerned with the impact of blurring in contrast enhancement. If you want to count the Chapter 7 discussion of S/H, that means six full chapters are devoted to this broad topic. I believe you will agree later that lavishing such space on the technique was worth it. There is, however, a problem in evaluating its effects.

All these blurring methods alter the picture in ways that do not correspond to any feature of the human visual system. They introduce artifacts, intelligent ones hopefully, that fool the viewer into perceiving stronger transitions.

The trick, as ever, is not to go too far. The good news is that the tiresome words *too far* are more specific this time than usual. They mean that the artificiality has to be strong enough to intrude upon the viewer by drawing attention