Lifting Detail from Darkness

skapps_0276

J. P. Lewis
Disney TSL
3100 Thornton Ave
Burbank CA 91505

This application sketch describes a high quality method for separating detail from overall image region intensity. This *intensity-detail* decomposition can be used to automate some specialized image alteration tasks. Our work was motivated by the movie *102 Dalmatians*, which featured a dalmatian puppy without spots. Animal handlers determined that the obvious idea of applying makeup to the dog was not possible – evidently there was no suitable makeup that was both safe to the dogs and that would stay in place during normal dog activity (including licking). This left Disney TSL with the task of painting out all the spots on the dog every few frames (the paintings were carried for a few frames with a simple scheme).

The spot removal task was larger than anyone guessed, and ultimately required a large number of artists (up to 40) working for eight months. The problem also proved to be more difficult than expected from an algorithmic point of view. As the spots often had visible fur texture, initially it was believed that there must be some simple compositing technique that could lighten these spots.

To get a feel for the problem, consider one representative approach inspired by unsharp masking: blur the dog and then subtract the blurred version from the original, giving a high-pass texture containing the fur. Then correct the intensity of the blurred dog to remove the spots (without saying how this is done). Lastly, add the high-pass to the lightened blurred dog, resulting (we hoped) in a dalmatian without spots but nevertheless having fur derived from the original texture. The problem with such an approach is suggested in Fig. 1: the scale of the blur must be exactly matched to the spot profile or there will be an overshoot/undershoot around the spot edge. This is not achievable since the spot transition regions have markedly different widths even on opposite sides of a single spot – some more adaptive technique is needed.

Adaptive Filtering

The intensity-detail decomposition problem is reminiscent of the problem addressed by Wiener filtering: to separate signal from noise. Making use of this observation, by casting detail in the role of ‘noise’ and intensity in the role of ‘signal’ we were able to apply a Wiener separation approach; a simple spatial-domain adaptive Wiener filtering algorithm described in [1] works quite well.

Intensity Modification

Once the detail is successfully separated we need a means of altering the image region intensity in a simple and controllable fashion. The membrane (Laplace) equation $\nabla^2 u = 0$ produces an interpolation of specified boundary conditions that minimizes $\int (\nabla u)^2 dA$ (the integrated gradient); as such it provides an adequate way of interpolating intensity specified on a boundary curve (e.g. a rough spot boundary). The Laplace equation is a linear system $Au = b$ with A being a sparse square matrix whose dimension is the number of pixels in the region. The intensity-detail decomposition was initially prototyped in MATLAB (which took less than a day) and used sparse matrix routines in that package. Our algorithm was later reimplemented in Java, which took about three months, approximately half of that being devoted to the membrane solution. The first implementation was a direct (non-sparse) matrix solution, programmed by Yi-Xiong Zhou. This was adequate for very small regions but was too inefficient for larger regions – areas of 150^2 pixels were requiring several minutes and 0.5G of memory! Fortunately the Laplace equation can be solved with the multigrid technique [2] (and in fact is the model problem for this approach). A multigrid implementation of the membrane reduced the solve time to several seconds even for large regions.

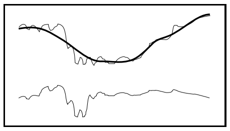

Figure 1: Hypothetical spot luminance profile, blurred luminance (heavy), and unsharp mask (bottom). When the blur size is too large the texture overshoots; when it is too small the blurred curve follows the texture and texture amplitude is reduced.

Figure 2: Demi-dalmatian.

Applications

In addition to the spot removal application, the intensity-detail decomposition has other specialized applications such as altering or removing shadows and reducing wrinkles (Fig. 3).

It should be emphasized that although the effects shown here are routine work for a photoshop artist,

- The image alterations shown were produced with no artistic skill in a few seconds each by providing only a crude outline for the desired region. Consider Fig. 3: the altered regions have luminance gradients as well as recovered original texture, so the effects could not be produced with a simple cloning operation but would require careful airbrushing followed by detail painting.
- Unlike manual retouching, the detail decomposition can be keyframed and produces consistent effects across frames.

References

[1] J. S. Lim, *Two-Dimensional Signal and Image Processing*, Prentice-Hall, 1990.

[2] W. H. Press, S. A. Teukolsky, W. T. Vetterling, B. P. Flannery, *Numerical Recipies*, Cambridge, 1993.

Figure 3: Altered regions include shadows under nose, eyes, and other changes.

Lifting Detail from Darkness

J. P. Lewis
Disney TSL
3100 Thornton Ave
Burbank CA 91505

skapps_0276

Supplementary material

Figure 4: Original dalmatian picture

Figure 5: Missing spots detail.

Lifting Detail from Darkness

J. P. Lewis
Disney TSL
3100 Thornton Ave
Burbank CA 91505

skapps_0276

Presentation Requirements

If the sketch is accepted I would like to demonstrate the software 'live'. For this I would need a machine of speed 400mhz or more (any OS) with 128M of memory. I can bring a laptop with the software preinstalled if it can be attached to the video projectors.