Supplementary figures

This technical report contains supplementary figures supporting the Eurographics 2006 paper

> "Cross dissolve without cross fade: preserving contrast, color and salience in image compositing", M. Grundland, R. Vohra, G. P. Williams, and N. A. Dodgson, *Computer Graphics Forum*, **25**(3), E. Gröller and L. Szirmay-Kalos (guest editors), 2006

Unless otherwise noted, the default parameters were used to produce all images shown.

Comparing standard linear blending with salience pyramid

The following images demonstrate how salience preserving pyramid compositing can produce pictures that are more visually appealing than standard linear compositing. To show the efficacy of our techniques, these examples were all produced with flat opacity maps. The alpha values are shown in square brackets below the source images.

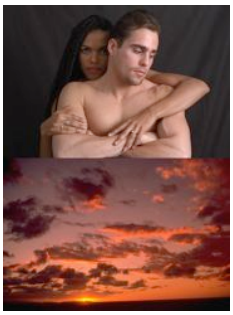

[0.6, 0.4]

Linear

Salience Pyramid
(custom parameters)

Compare salience pyramid compositing to the double exposure effect of linear compositing. Note, in particular, the opacity of the man's chest and the contrast retained by the sky and the man's hair. Notice how the woman's face blends with the sky to create a more æthereal composite. Creating such an effect such would be time consuming using traditional methods.

[0.55, 0.45]

Linear

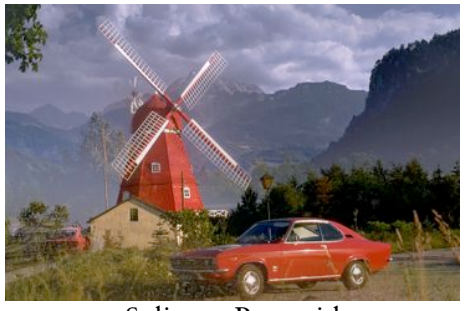

Salience Pyramid

Notice how salience pyramid compositing has automatically boosted the weighting given to the windmill and the cars in both components: they are more opaque than in the linear composite. In particular, observe how well the sails stand out from the background and that the sky has far better contrast.

[0.5, 0.5]

Linear

Salience Pyramid

In these images, notice how salience preserving pyramid compositing has made the birds more opaque than the linear composite. In addition the contrast of the forest is substantially higher. The salience preserving operator has not completely cut the birds from their original background, as evidenced by the blue mist covering the forest. However, in this composite, the blue mist adds to the ambiance of the scene. One remaining problem is that the wings of the topmost bird are not completely opaque. A user drawn opacity map could easily alleviate this.

[0.8, 0.2]

Linear

Salience Pyramid

Salience pyramid compositing renders the skiers as solid figures. However, notice that the blending of the upper skier with the rock formation is not entirely convincing. These imperfections could be ameliorated by a simple user supplied opacity map.

[0.75, 0.25]

Linear

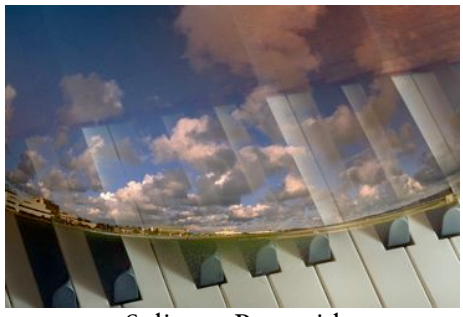

Salience Pyramid

This composite shows far greater contrast when combined with our salience pyramid operator than with linear averaging. In particular notice how both the clouds and the piano keys remain prominent in the salience preserving composite.

[0.5, 0.5]

Linear

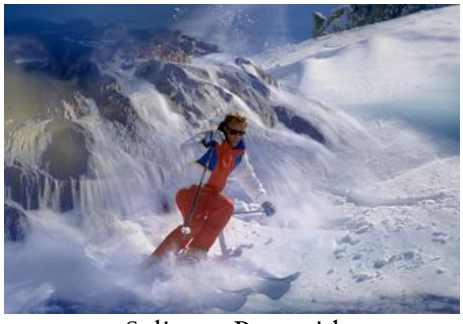

Salience Pyramid

Notice the greater contrast, and opacity of the foreground and background elements in the salience pyramid composite. This picture is compositionally interesting because the picture of the wave breaking over the rocks is interpreted as water in isolation, but as snow or ice when blended with the skier.

Salience mattes

These images show how mattes, used in salience preserving compositing, are automatically derived from the salience maps. Studying these reveals how salience preserving compositing generates the set of salience mattes with which to combine the images. The automatically generated mattes would typically have been very time consuming to produce with current techniques.

Note that these figures just show the salience maps and salience mattes used at the highest resolution of multiresolution salience preserving pyramid blending.

Original Images

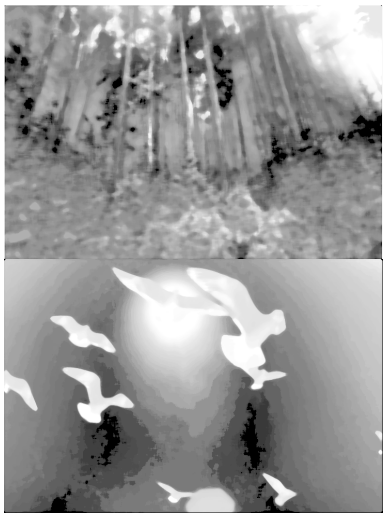

Salience Maps

Salience Mattes: α=[0.5, 0.5]

In this example, we can see that the transparency in the topmost bird's wings is owing to the coincidence of high saliency between the wing in one image and the light shining through the trees in the other.

Original Images

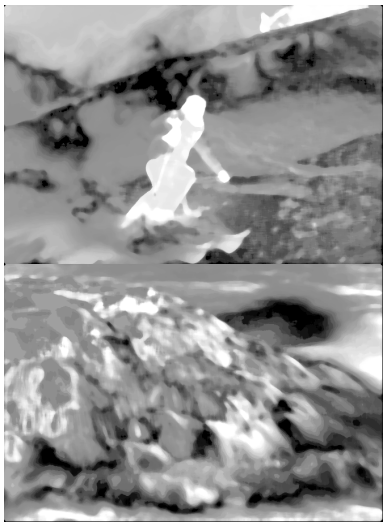

Salience Maps

Salience Mattes: α=[0.5, 0.5]

Our saliency maps are exclusively based on colour features. They thus easily pick out the coloured parts of the skier's jacket, but not necessarily the entire skier. In particular the white parts of the skier's clothing are the same colour, and therefore the same saliency, as the snow.

Three-way compositing and compositing operator comparisons

These sets of images each show three images being composited together using all four of the compositing operators. Flat opacity maps are used in each case. Although subtle in some cases, these examples demonstrate the visual differences between our three new operators and the standard linear composite.

[0.33, 0.33, 0.33]

Linear

Contrast

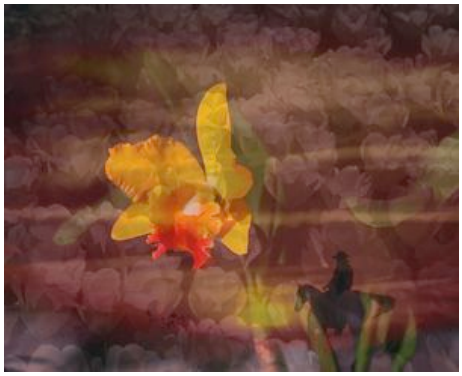

Colour

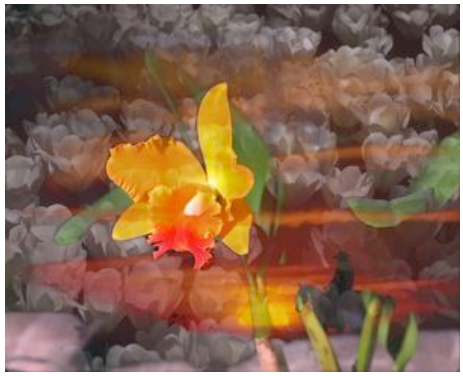

Salience

Our contrast preserving method retains the high contrast of the white flowers much better than linear compositing. Our colour preserving operator maintains stronger tones of red on the background and of yellow in the flower. Our salience preserving operator clearly brings the yellow flower to the fore, whilst also making the area surrounding the horse, rider and setting sun less cluttered, and it has given prominence to the white flowers in the upper right corner, where the other two images have low salience.

[0.4, 0.2, 0.4]

Linear

Contrast Pyramid

Colour Pyramid

Salience Pyramid

Our contrast pyramid method demonstrates how multiple overlaying images can remain recognisable at a glance when combined in a visually coherent composition. Our colour pyramid method is effective at bringing out the strong bold colours of the components, as visible in the front of the barn, the rocks and the reeds. Our saliency pyramid method gives emphasis to particular aspects of the component images, in this example the sunset is particularly prominent compared with the other three methods.

Non-uniform masks

The following image sets show how our operators behave when given non-constant opacity maps. They also show how the opacity maps and saliency maps are combined to produce salience mattes.

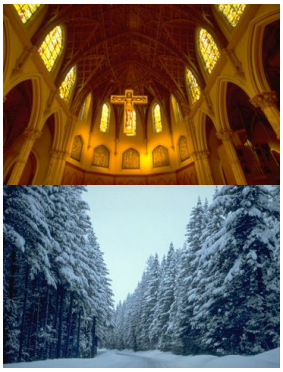

Source Images

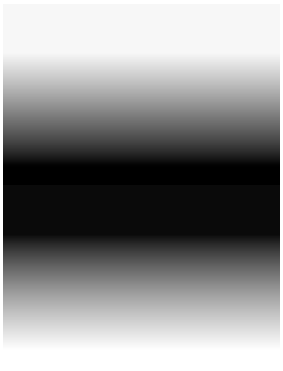

Opacity Maps

Salience Maps

Salience Mattes

Linear Pyramid

Contrast Pyramid

Equally Weighed Linear Average of Contrast, Colour and Salience Composites

Colour Pyramid

Salience Pyramid

In this example, the differences between the methods can be clearly seen. Our contrast preserving method has noticeably more contrast than linear compositing. Our colour preserving composite maintains a better yellow tone. Our salience preserving method highlights the lower arches of the church and the tops of the trees, features which are not weighted highly by the opacity maps. Finally, the linear combination of our three methods shows some aspects of each method and demonstrates that additional artistic effects can be created simply by combining the different compositing methods.


Source Images

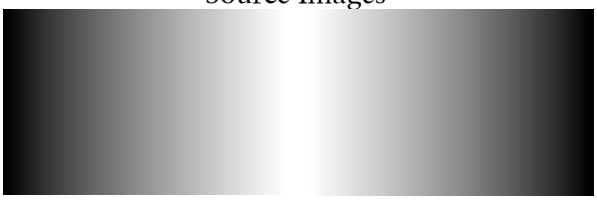

Opacity Maps

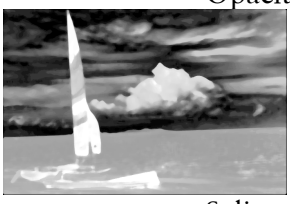
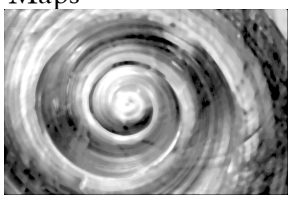

Salience Maps

Salience Mattes

Linear

Salience


Source Images

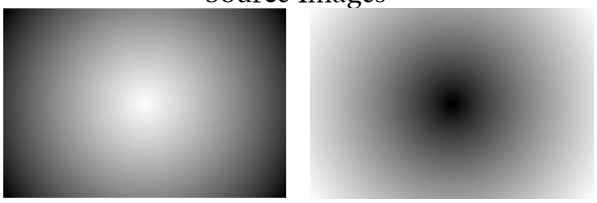

Opacity Maps


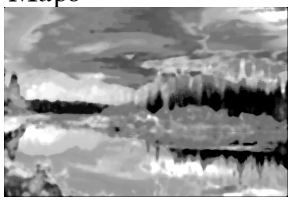

Salience Maps

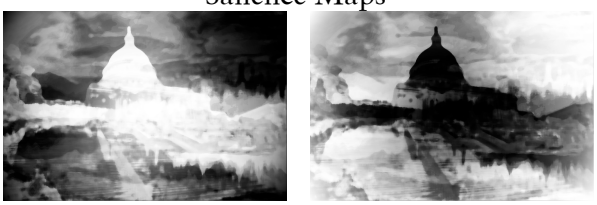

Salience Mattes

Linear Pyramid

Salience Pyramid

Source Images

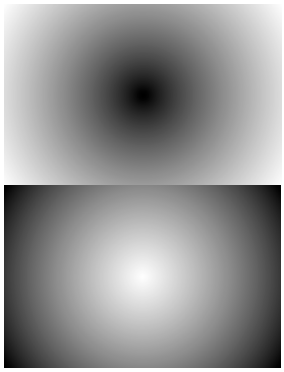

Opacity Maps

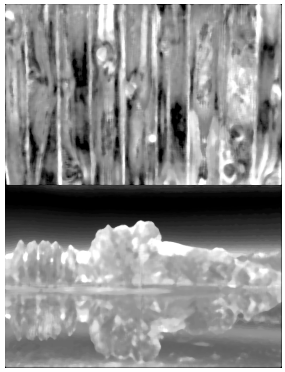

Salience Maps

Saliency Mattes

Linear Pyramid

Contrast Pyramid

Colour Pyramid

Pyramid Saliency

User-drawn opacity maps

The following sets of images show how the user can specify a simple opacity map to aid the compositing operators. In the case of salience preserving compositing, the fine details can then be derived from the saliency maps.


Source Images


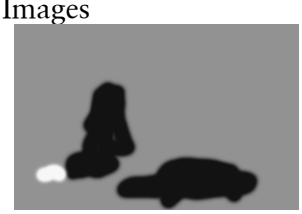

Opacity Maps


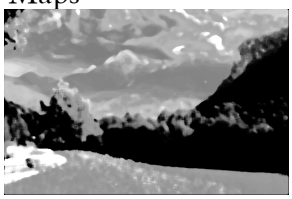

Salience Maps

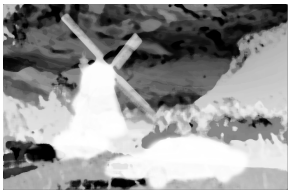
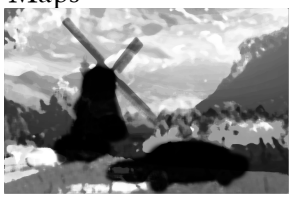

Salience Mattes

Linear Pyramid

Salience Pyramid

In this example notice how salience pyramid compositing has automatically maintained the bright white windmill sails even though the user-drawn opacity map does not include them. Also, notice the fringe around the windmill in the linear composite caused by the lack of precision in the outline of the opacity map. This effect is reduced in the salience pyramid composite because of the multi-resolution behaviour of the operator.


Source Images




Salience Maps

Opacity Maps


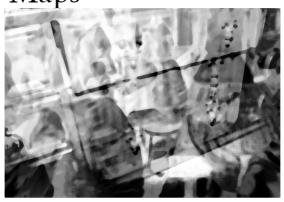

Salience Mattes (using flat opacity maps, alpha=0.5)

Salience Mattes (using user drawn opacity maps)


Linear Pyramid

Linear Pyramid


Salience Pyramid
The above two images use flat opacity maps

Salience Pyramid
The above two images use the user's opacity maps

This example compares the same composite using both a user supplied opacity map and a flat opacity map. When using the flat map the characters' heads are missing because the relative salience of the cabinet is greater than that of the heads. This problem can be corrected by quickly drawing an opacity map to force the inclusion of the heads.

Cross dissolves

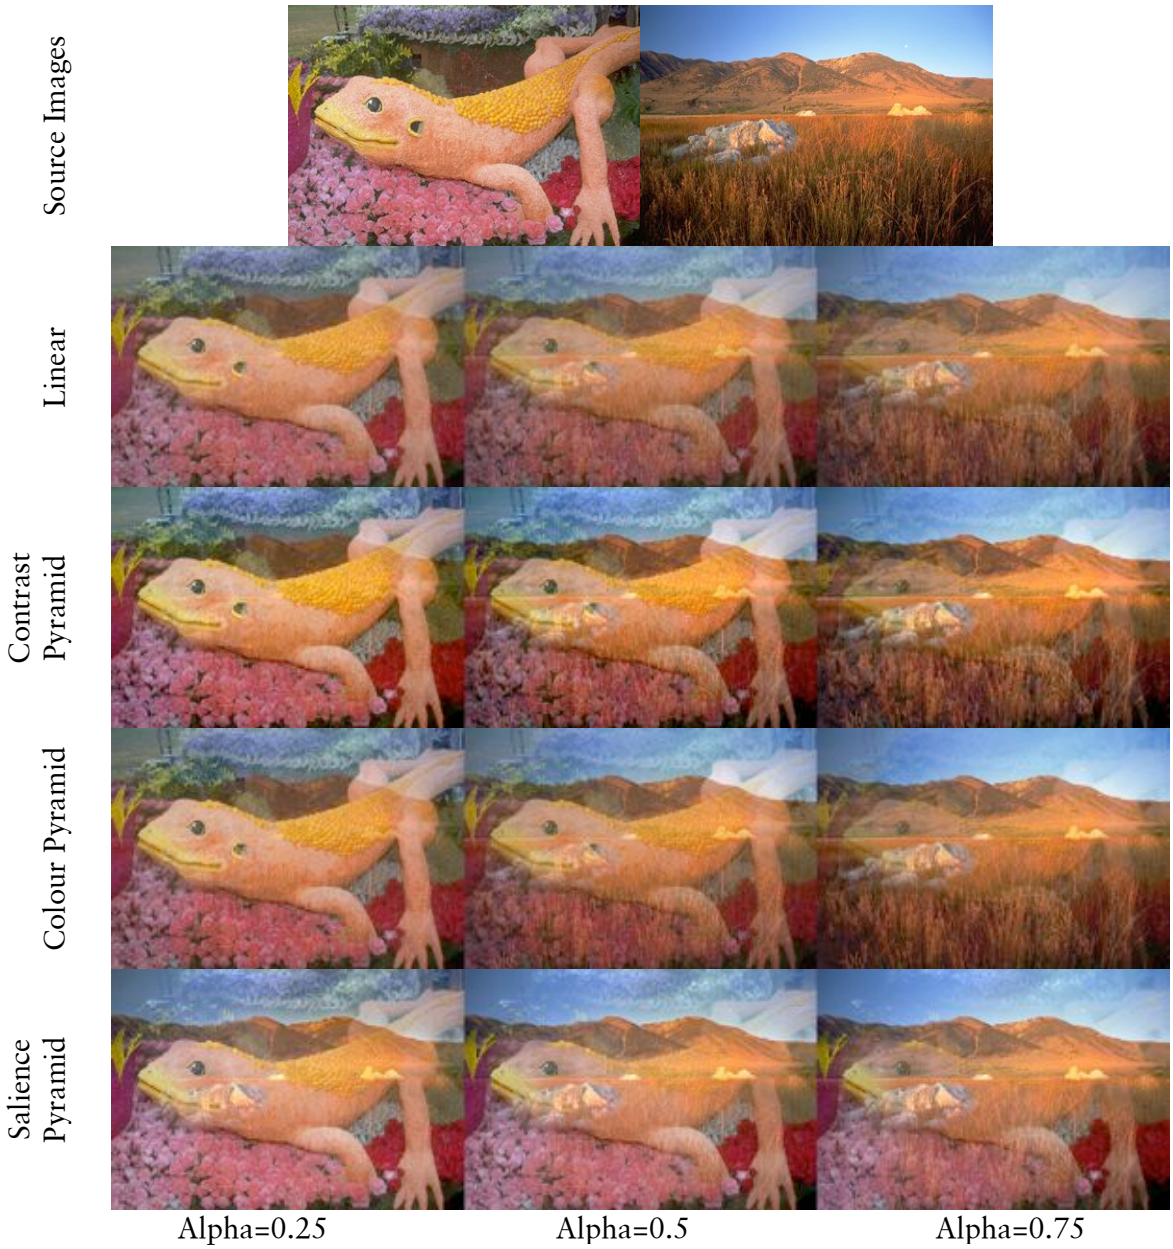

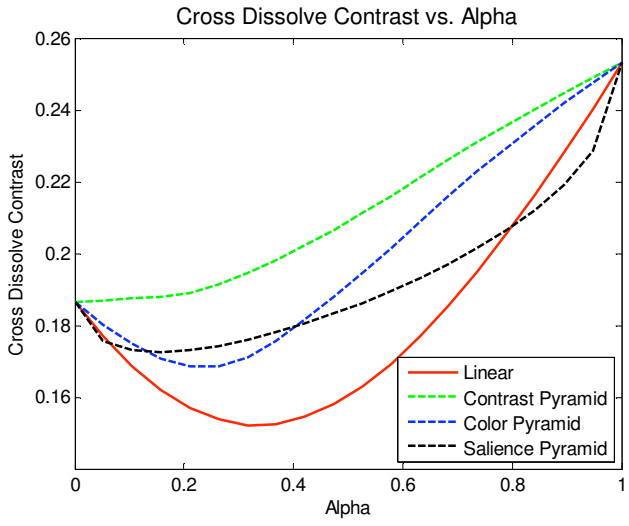

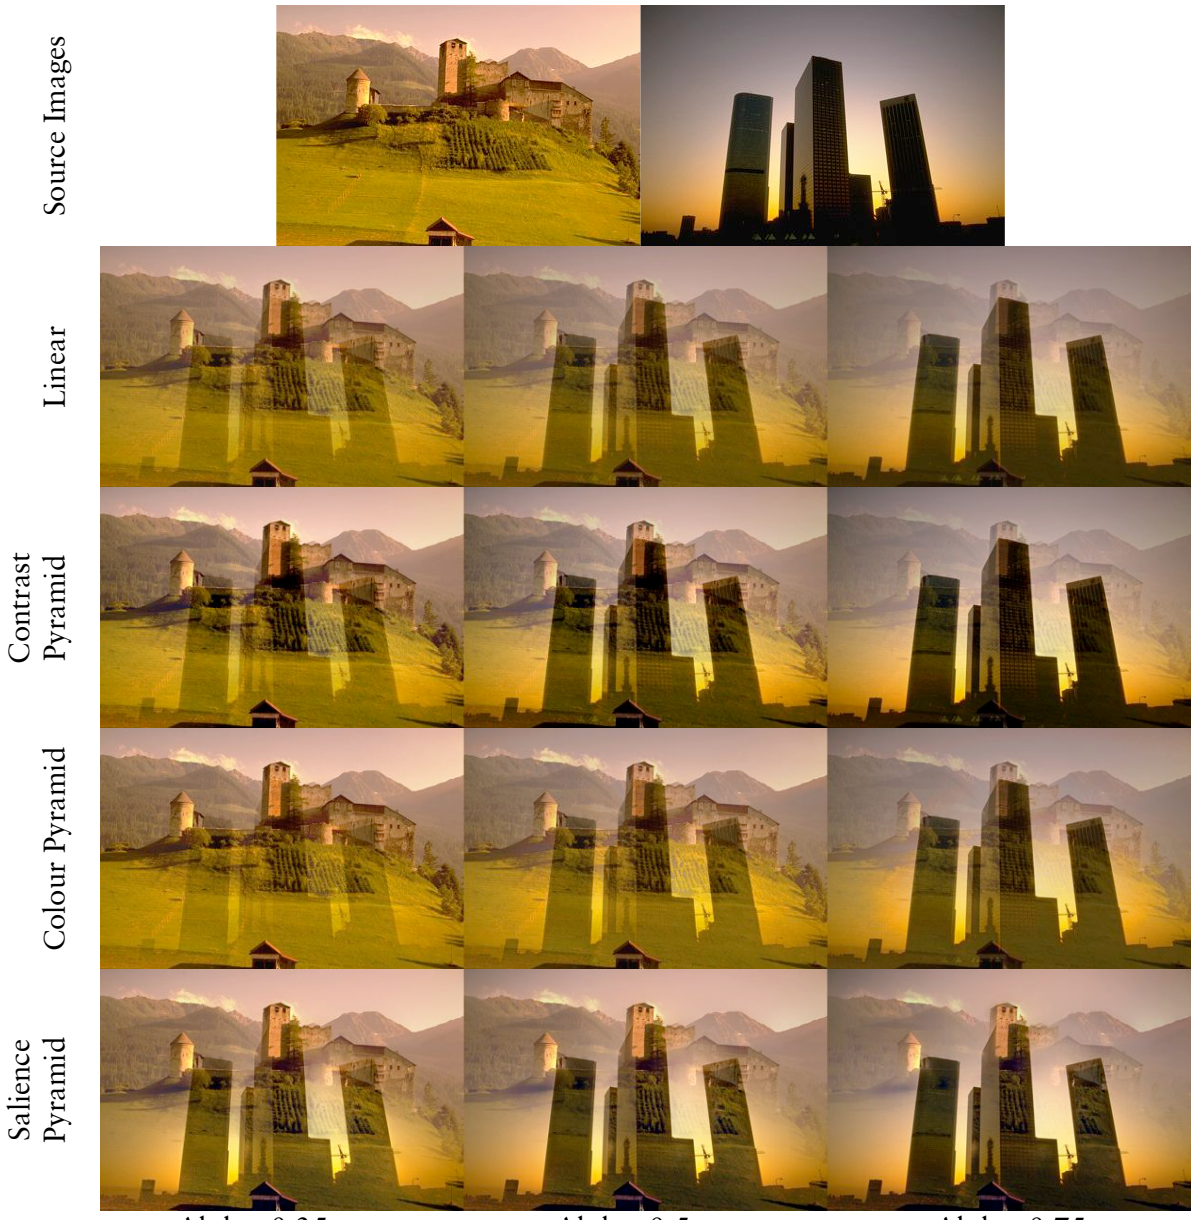

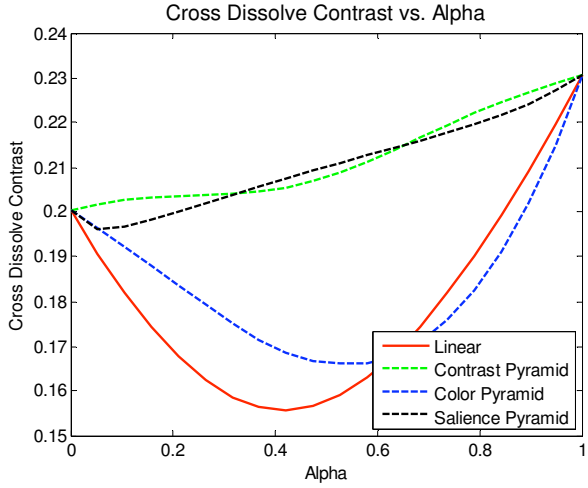

In this example, note that although the contrast of the colour preserving blend is low, it is still visually appealing because vibrant colours are maintained.

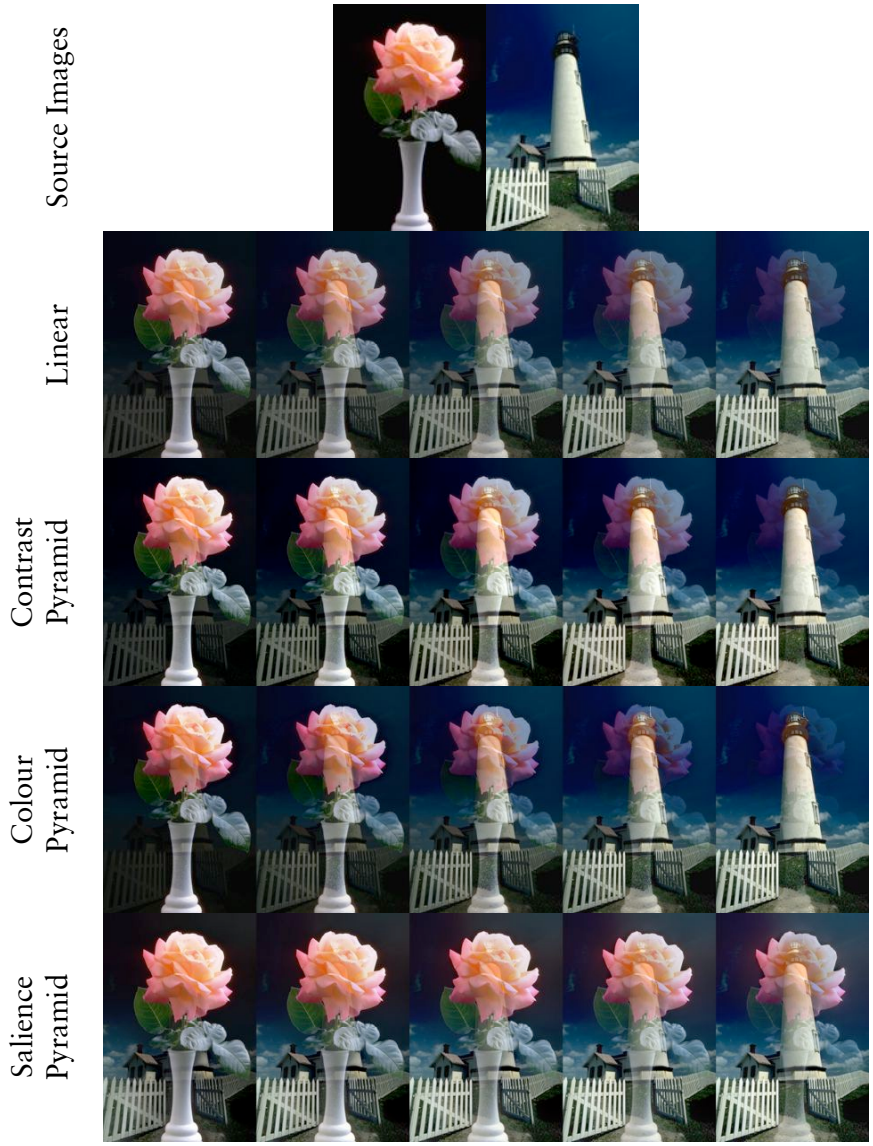

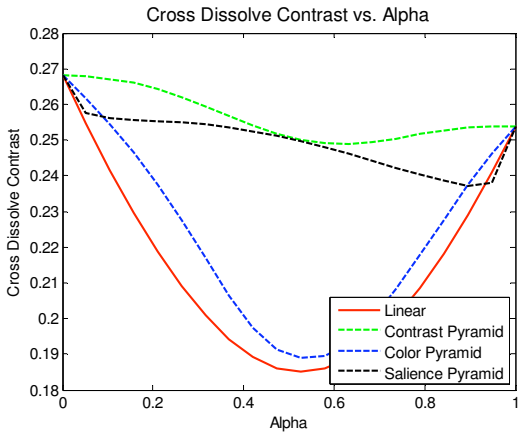

Observe that in this example, our contrast pyramid method doesn't exactly preserve contrast at the mid point of the blend. This can be seen in the dip of the graph of its contrast level. This artefact is caused by the fact that contrast preserving and contrast pyramid blending can, on occasion, map a small portion of colours out of gamut, and thus its contrast enhancing effect is clipped. This example also produces an interesting salience pyramid blend where the picket fence enters first and then the sky. Finally, the vase turns into the lighthouse with the petals the last to disappear.

Filtering

Linear

Contrast

Colour: $\lambda = e^{-8}$

Linear

Contrast

Colour: $\lambda = e^{-8}$

Contrast and colour preserving image blending can be adapted to image filtering. To apply contrast preserving filtering, we filter the image normally and then linearly stretch each colour channel about its mean to regain the contrast of the input. To apply colour preserving filtering, we map colours to real values, filter these, and then map back to colour values. These examples show the results of Gaussian blurring using linear filtering, contrast preserving filtering and colour preserving filtering with $\lambda = e^{-8}$. The latter produces an interesting blur which preserves high contrast edges.

Inside/Outside

The following images show how salience pyramid blending can be used to show the interior and exterior of objects simultaneously.

[0.95, 0.05]

Linear

Salience Pyramid

In this example, blending 95% of the computer exterior with 5% of the computer interior produces a boring image with linear blending, but salience pyramid blending produces a translucent case showing vivid power cables. The power cables shine through because they have high colour entropy.

[0.75, 0.25]

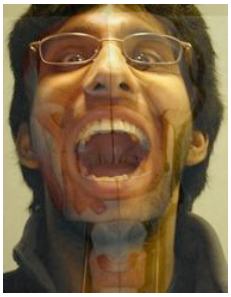

Linear

Salience Pyramid

In this example, notice how salience pyramid blending shows the vertebra and musculature through the neck and the left side of the face. These parts shine through because they have high colour entropy.

Stroboscopic Visualisation of Motion

[0.33, 0.33, 0.33]

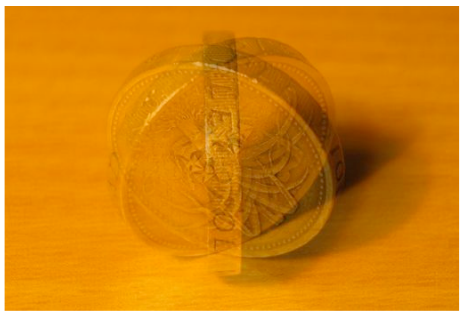

Linear

Salience Pyramid

In this example, notice how salience preserving pyramid blending retains the details along the edge and on the face of the coin, producing a solid looking result.

Compositing Image With Its Blur

Original

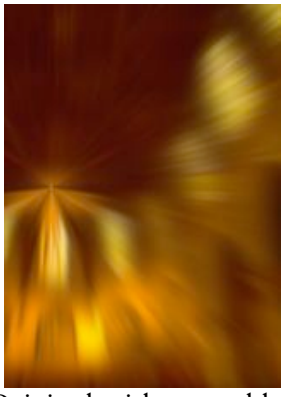

Original with zoom blur

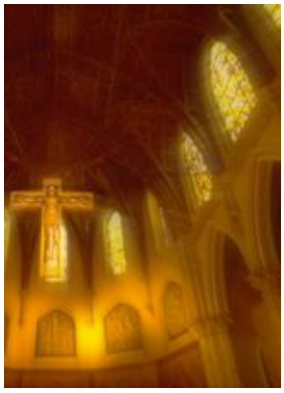

Linear [0.4, 0.6]

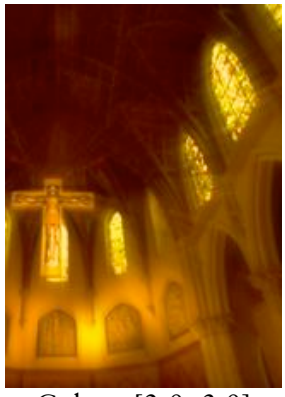

Colour [2.0, 3.0]

This example shows colour preserving blending with non-convex weights, producing much sharper and brighter windows while preserving the subtle zoom blur.

Mipmapping

Trilinear filtering is a common way to produce mipmaps. It uses linear interpolation both to produce the mipmap levels and also to interpolate between them. This causes texture contrast, defined as the mean standard deviation of the RGB colour channels, to dip sharply between mipmap levels. The following figures show the results of using contrast preserving blending to linearly interpolate texture contrast between mipmap levels. This approach avoids sudden changes in contrast as magnification varies across mipmap levels.

Wicker

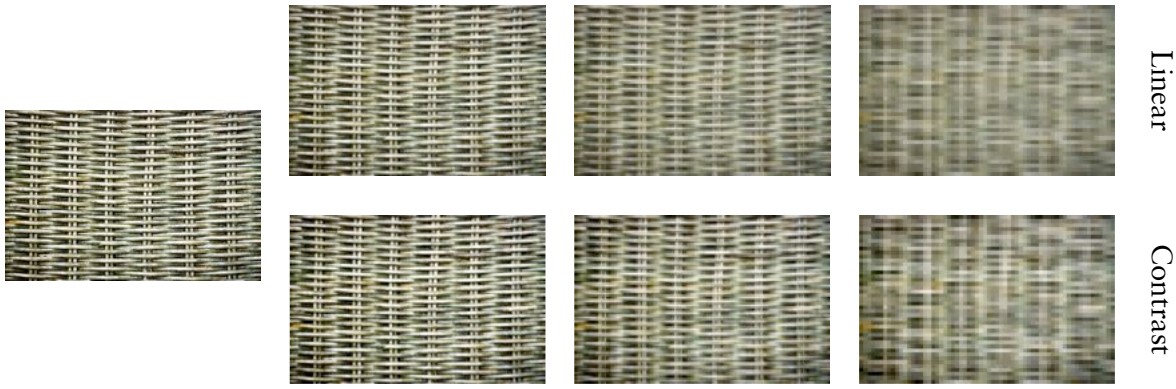

Original Texture | M = 0.75 | M = 0.375 | M = 0.1875

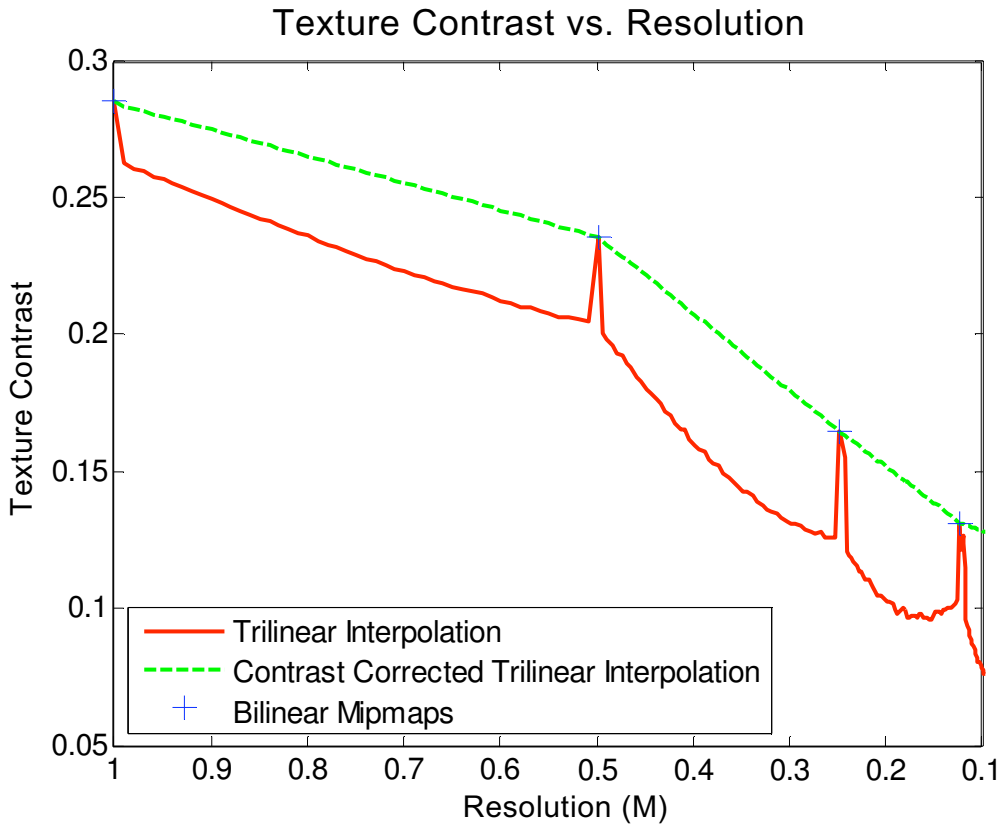

Hessian

Original Texture M = 0.75 M = 0.375 M = 0.1875

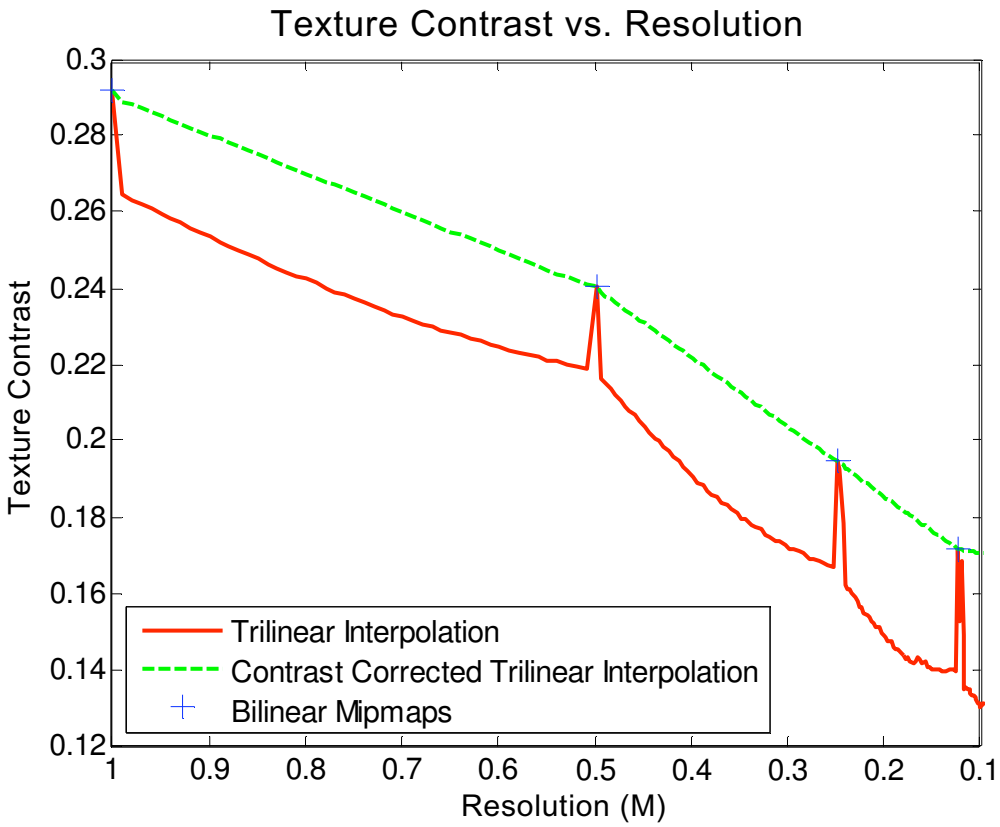

Pebbles

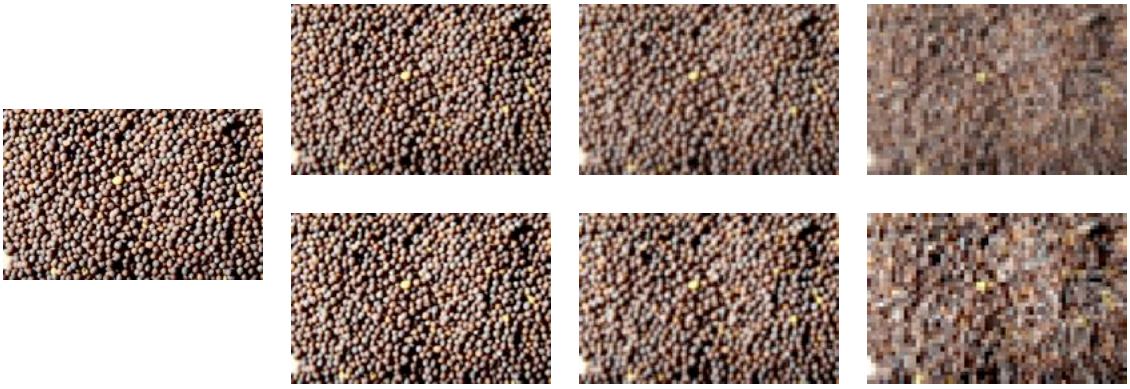

Original Texture M = 0.75 M = 0.375 M = 0.1875

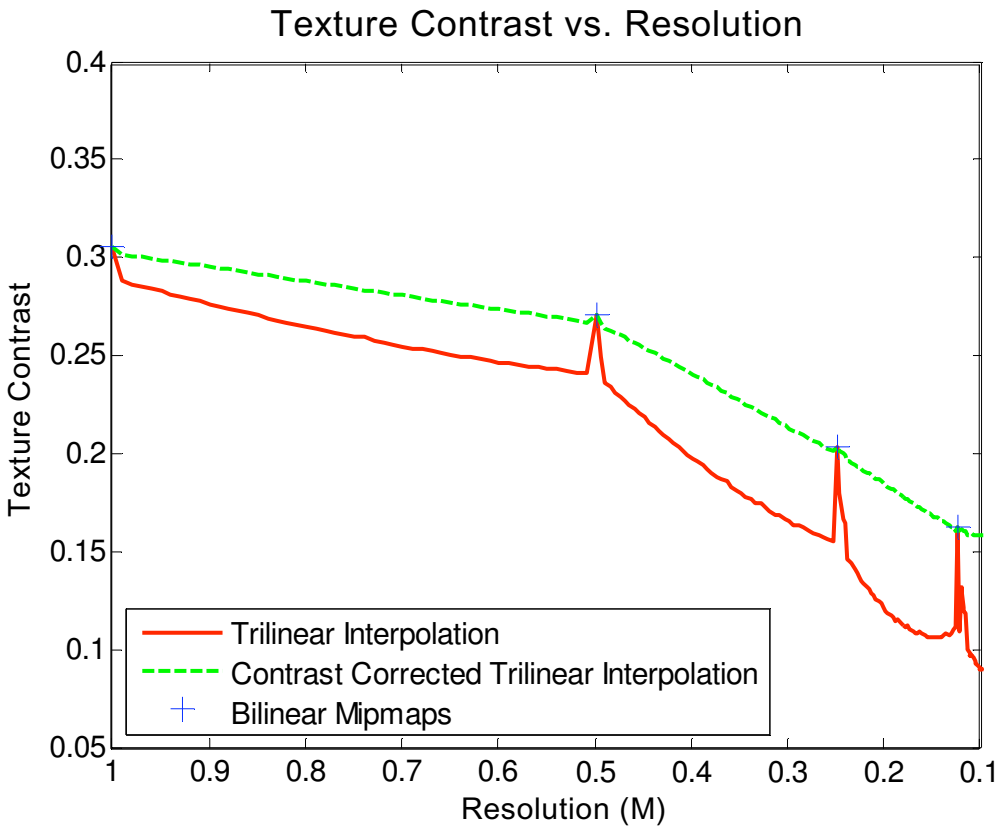

Image Morphing

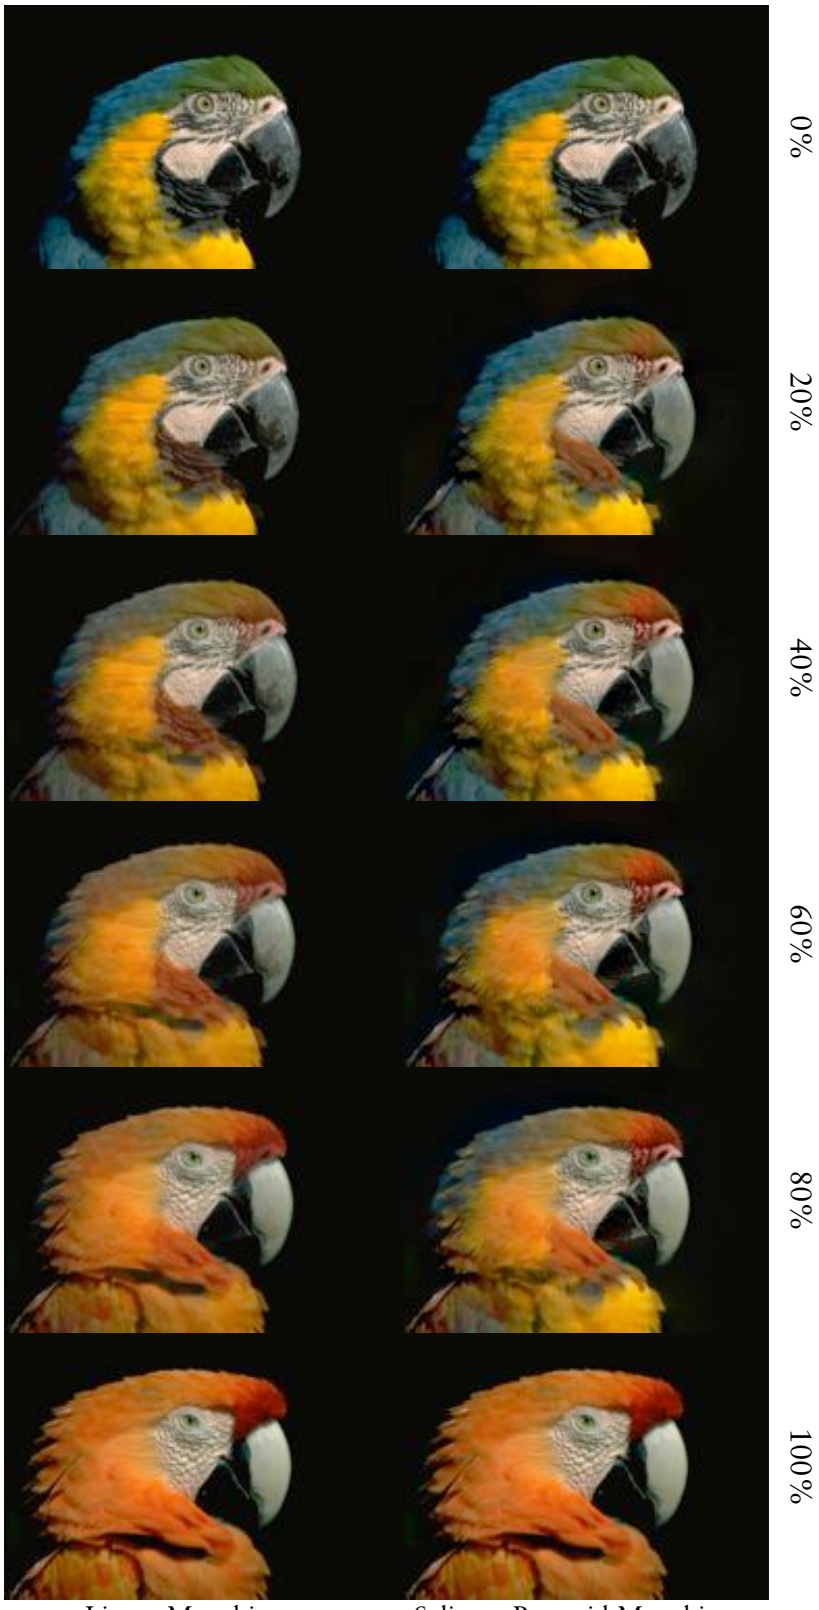

Linear Morphing Salience Pyramid Morphing

Image morphing linearly averages two geometrically warped image sequences, causing contrast and ghosting artefacts in areas of dissimilar texture and colour. This example shows how salience pyramid blending can be used to avoid these artefacts. Notice that at 40% and 60%, salience preserving keeps the blue feathers distinguishable at the top of the head, whereas linear blending simply blends to an intermediate grey.