

# Image compression using sparse colour sampling combined with non-linear image processing

Stephen Brooks<sup>\*a</sup>, Ian Saunders<sup>b</sup>, Neil A. Dodgson<sup>\*c</sup>

<sup>a</sup>Dalhousie University, Halifax, Nova Scotia, Canada B3H 1W5

<sup>b</sup>University of Edinburgh, Edinburgh, Scotland EH1 2QL

<sup>c</sup>University of Cambridge, Cambridge, England CB3 0FD

## ABSTRACT

We apply two recent non-linear, image-processing algorithms to colour image compression. The two algorithms are colorization and joint bilateral filtering. Neither algorithm was designed for image compression. Our investigations were to ascertain whether their mechanisms could be used to improve the image compression rate for the same level of visual quality. Both show interesting behaviour, with the second showing a visible improvement in visual quality, over JPEG, at the same compression rate. In both cases, we store luminance as a standard, lossily compressed, greyscale image and store colour at a very low sampling rate. Each of the non-linear algorithms then uses the information from the luminance channel to determine how to propagate the colour information appropriately to reconstruct a full colour image.

**Keywords:** compression, sparse, non-linear, colour, colorization, bilateral filter

## 1. INTRODUCTION

It is widely known that the human eye is far more responsive to luminance than to chrominance<sup>1</sup>. Recent research has investigated ways of automating the process of “colorization”: adding colour to monochromatic content, such as black & white movies<sup>2,3,4</sup>. Our research investigated the combination of the two: if we sample chrominance at low resolution, can these colorization algorithms recover a sufficiently good rendition of the image to be useful in colour image compression.

The two algorithms are one explicitly named “colorization” by its creators<sup>4</sup> and the joint bilateral filter<sup>5</sup>. Neither algorithm was designed for image compression. Our investigations were to ascertain whether their mechanisms could be used to improve compression rate for the same level of visual quality. Both show interesting behaviour, with the second showing a visible improvement in visual quality, over JPEG, at the same compression rate.

In both cases, we store luminance as a standard, JPEG compressed, greyscale image and store colour at a very low sampling rate. Each of the non-linear algorithms then uses the information from the luminance image to determine how to propagate the colour information appropriately to reconstruct a full colour image.

Colorization<sup>4</sup> is a method developed to convert a greyscale image to colour using a minimal amount of user intervention. The user specifies the colour at a relatively small set of locations in the image. The algorithm then propagates colour information from these locations under the assumption that adjacent pixels with similar luminance are likely to have a similar colour. In our experiments, we sub-sampled the colour information and then fed those colour samples into the algorithm as single colour pixels regularly spaced in a sea of greyscale pixels.

The joint bilateral filter<sup>5</sup> is a mechanism whereby two images, of the same scene, are combined to produce an improved final image. It is used in a range of applications including “flash/no-flash” image processing<sup>6,7</sup>, in which the natural illumination of the scene is inadequate to provide a crisp image with a short exposure time. One of the images is captured using a flash, the other with no flash. The “flash” image will tend to have good colour and detail but these are obtained by sacrificing subtle shadows, reflective interactions between objects, and the natural lighting of the scene. The “no flash” image will contain the subtle shadows, the natural lighting, and more “moody” colour. However, the “no flash” image will tend to be very noisy, owing to inadequate illumination. In our application, since the luminance

---

\* Stephen Brooks sbrooks@cs.dal.ca, Neil Dodgson nad@cl.cam.ac.uk

channel has good edge detail, the luminance channel can be thought of as analogous to the “flash” image. The low-resolution colour channels can be reconstructed, by nearest-neighbour sampling, to produce a blocky colour image, which can be treated as the “no flash” image. When these images are fused, using a joint bilateral filter, the blockiness of the colour channel is spread out to match more closely the edges within the luminance channel.


In the case of colorization, we compared our results against standard JPEG compression. In our experiments, we sampled the colour information at a range of spacings, from every second pixel to every nineteenth pixel (the latter thereby reducing the colour information by a factor of 361). We combined this colour information with a standard JPEG compressed greyscale image at a range of JPEG compression rates. Our experiments showed that this method introduced a range of non-standard artefacts different to those introduced by JPEG colour compression. In particular, for the same bit rate as JPEG, it tends to have fewer blocky artefacts but more washed out colour. We compared PSNR values between JPEG compressed imagery and images compressed by colorization. We considered what sample spacing, in compression by colorization, matched to what JPEG compression number from the IJG implementation of JPEG ([www.ijg.org](http://www.ijg.org)). Different types of imagery exhibited different characteristics. Many images exhibited a linear relationship. Some images, notably those with large smooth areas of colour, exhibited better performance under compression by colorization than they did under JPEG. In particular, they had very good performance up to a sub-sampling rate of one colour sample for every  $6\times 6$  luminance pixels. Other images, notably those with colour that varied artificially quickly (for example, Figure 9), performed worse under compression by colorization than under JPEG. In addition, our experiments show that the degradation of image quality with respect to colour sample spacing, in compression by colorization, is not uniform across all images nor is it always monotonic. Despite the ambivalent results, this is an interesting first look at using the colorization algorithm in image compression and it offers a starting point for exploring these atypical approaches to image compression.

The joint bilateral filter algorithm performs somewhat more consistently. In our experiments on this method, we compared standard compression of the colour image, using JPEG, with compression of the colour information at a very low rate combined with compression of the greyscale information using JPEG at a rate such that the overall number of bits stored was equivalent to that in the standard compression. The joint bilateral filter algorithm gave an improved visual result over the standard JPEG method; the most important feature being the dramatic reduction in spurious colour shift artefacts. There is of course a penalty in the time required to run the joint bilateral algorithm. Our fast implementation requires less than a minute for a  $1280\times 1024$  image, and increases in computation power will increasingly make methods like this feasible. Overall, this method offers improved quality for the same bit rate at the expense of increased processing time in decompression.

## 2. COMPRESSION BY COLORIZATION

We now discuss the first approach in more detail. We begin with the observation that, in many images, there is a great deal of colour coherence. In particular, most images consist mainly of regions of smoothly varying colour. This suggests that we can store colours at a subset of locations and subsequently generate the necessary gradients through a process of optimization. The recent work on colorization<sup>4</sup> offers a starting point for exploring this atypical approach to image compression. Image colorization is a method developed to convert a greyscale image to colour using a minimal amount of user intervention. The user specifies the colour at a relatively small set of locations in the image. The algorithm then propagates colour information from these locations under the assumption that adjacent pixels with similar luminance are likely to have a similar colour.

In our experiments, we store luminance as a standard, compressed, greyscale image and store colour at a low sampling rate. The colour samples are regularly spaced in a sea of greyscale pixels. Figure 1 shows an example of this, with the left side of the figure displaying a zoomed region of an image. If one looks closely, colour values are retained only sparsely on a grid. Specifically, our method of compression retains the luminance values (L) computed in  $L\alpha\beta$  colour space<sup>5</sup>, which is designed for perceptual uniformity. We retain the subset of  $\alpha\beta$  colour component values at regular grid spacing, which we name Chrominance Points (CPs).



**Fig. 1.** We store the greyscale image and a set of chrominance points (left). This allows us to generate an approximate reconstruction (right) of the original image (centre). This example is at a high level of colour compression

To approximately reconstruct the image, the CPs are fed into the image colorization algorithm along with the greyscale image. In our initial work we have performed experiments with the existing colorization method of Levin *et al.*<sup>4</sup> which minimizes the difference between the colour  $\alpha\beta$  values  $\alpha(x, y)$ ,  $\beta(x, y)$  at pixel  $(x, y)$  and the weighted average of the colour values at neighbouring pixels,  $(x', y')$ . This is based on the assumption that neighbouring pixels will likely have similar colours if they have similar intensities. For each chrominance channel,  $\alpha$  and  $\beta$ , the following is minimized:

$$\sum_{(x,y)} \left( \alpha(x,y) - \sum_{(x',y') \in N(x,y)} w(x,y,x',y') \alpha(x',y') \right)^2 \quad (1)$$

where  $w(x, y, x', y')$  is a weighting function summing to one, which is large when intensity levels at pixel locations  $(x, y)$  and  $(x', y')$  are similar. Specifically, the weighting function is based on the normalized correlation between the two intensities:


$$w(x, y, x', y') \propto 1 + \frac{1}{\sigma_{x,y}^2} (L(x, y) - \mu_{x,y}) (L(x', y') - \mu_{x,y}) \quad (2)$$

where  $L(x, y)$  is the intensity at pixel  $(x, y)$ , and where  $\mu_{x,y}$  and  $\sigma_{x,y}$  are the mean and variance of the intensities in a window around  $(x, y)$ . In this way, the algorithm is able to use information from the luminance image to determine how to propagate the colour information appropriately to reconstruct a full colour image.

### 3. RESULTS FOR COMPRESSION BY COLORIZATION

For colorization-based image compression, we compare our results against standard JPEG compression. In our experiments, we sampled the colour information at a range of spacings, from every second pixel to every nineteenth pixel (the latter thereby reducing the colour information by a factor of 361). We combined this colour information with a standard JPEG compressed greyscale image at a range of JPEG compression rates. The greyscale values are compressed as a single channel JPEG image, while the  $\alpha\beta$  colour component values are scaled and stored separately in a green-blue losslessly compressed PNG file. We experimented with a number of lossless compression formats for storing  $\alpha\beta$ , including JPEG2000 (lossless), JPEG-LS, TIFF and PNG. PNG offered the best performance for our purposes.

Firstly, it is interesting to consider the nature of the compression artefacts our method produces and how they differ markedly from those produced with more traditional methods such as Discrete Cosine Transform (DCT) based approaches<sup>1,8</sup>. In particular, for the same bit rate as JPEG, our method tends to generate fewer blocky artefacts but at the expense of a more washed out colour. Figure 2 shows a comparison of typical artefacts at high levels of compression. While JPEG (centre) generates familiar wavelet-like artefacts; ours (right) loses colour fidelity and saturation, and at extreme levels of compression colours may bleed into adjacent areas. We speculate that the best way to mitigate these colour fidelity artefacts will be to adapt the placement of Chrominance Points (CPs) based on image content. This would, however, require us to store the locations on the CPs and therefore the advantage gained by having arbitrary locations would need to be more than offset the extra storage required.



**Fig. 2.** Examples at high levels of compression. Left: original. Centre: JPEG compression artefacts. Right: “compression by colorization” artefacts: the JPEG artefacts have practically vanished but there is some loss of vividness.


Secondly, our experiments show that the degradation of image quality with respect to grid spacing is not uniform across all images nor is it always monotonic. Image degradation is often dependant on image content. These characteristics emerged when we compared PSNR values between JPEG compressed imagery and images compressed by colorization. For this, we considered what sample spacing, in compression by colorization, matched to what JPEG compression number from the IJG implementation of JPEG ([www.ijg.org](http://www.ijg.org)). Figures 3–10 feature images that exhibit the various characteristics. In the associated graphs, equal PSNR values are plotted for JPEG compression (the ordinate is the JPEG compression number) and our method (the abscissa is the sample spacing).

Many images exhibited a linear relationship. Figures 5–8 show results with comparable image degradations to JPEG compression. In these cases, we note how there are no significant discontinuities in the curves.


Some images, notably those with large smooth areas of colour, exhibited better performance under compression by colorization than they did under JPEG. In particular, they had very good performance up to a sub-sampling rate of one colour sample for every 6×6 luminance pixels. Note how, in Figure 3, as grid spacing increases from 1 to 5, the PSNR result continues to equate to the PSNR produced by the maximum JPEG quality level. Degradation improvements are not as dramatic but still interesting on the Lena image that follows in Figure 4. The difference may be because Lena contains significant monochromatic fine detail in the feather and the hatband.

Other images, notably those with colour that varied artificially quickly, performed worse under compression by colorization than under JPEG. An example of such an image is shown in Figure 9. Here we see that the new method cannot produce good results even when colour samples are taken only every third pixel in both directions. This is probably owing to the extreme colour variations in this artificial image. In addition, our experiments show that the degradation of image quality with respect to colour sample spacing, in compression by colorization, is not always monotonic. Figure 10 shows an erratic curve for an image that contains large, smooth monochromatic areas. We speculate that these large monochromatic areas may be causing the unusual jaggedness of the graph.


Despite the ambivalent results, this is an interesting first look at using a colorization algorithm for the purpose of image compression and it offers a starting point for exploring these atypical approaches to image compression. Moreover, our experiments also point to the open question of how to best compare such different visual artefacts as those wavelet-like artefacts produced by JPEG and the reduced colour fidelity produced by ours. We know that PSNR only loosely correlates with human perception but, given the lack of any quantitative estimate of human perceptual quality, PSNR is the method in common usage. Just how well the PSNR correlates with perceived quality is open to further examination.




**Fig. 3.** For this image, our new method improves on JPEG for small compression factors. Here we see that the new method can preserve image quality better than JPEG, even when colour samples are taken only every fourth or fifth pixel in both directions.




**Fig. 4.** A less dramatic example of improved performance for the Lena image. This may be because Lena contains significant monochromatic fine detail in the feather and the hatband.




**Fig. 5.** For the sailing image, the degradation in quality, measured by PSNR, is similar for JPEG and our new method.




**Fig. 6.** For the mandrill image, the degradation in quality, measured by PSNR, is similar for JPEG and our new method.




**Fig. 7.** For the butterfly image, the degradation in quality, measured by PSNR, is similar for JPEG and our new method.



**Fig. 8.** For the pepper image, the degradation in quality, measured by PSNR, is similar for JPEG and our new method.



**Fig. 9.** An artificial image with rapidly changing colour patterns. The degradation in quality, measured by PSNR, is worse than for JPEG.



**Fig. 10.** For the milk drop image, the degradation in quality, measured by PSNR, is erratic.

**Note on figures:** Figures 3–10 contain original uncompressed versions of the image. All other images in this document should be viewed in colour on a monitor in order to evaluate visually the artefacts present in the compressed versions. You can access a PDF version of the paper from the symposium CD-ROM or from Dr Dodgson's website.

#### 4. JOINT BILATERAL FILTER

Our second method, the joint bilateral filter, was applied to images where both luminance and chrominance were compressed using JPEG. We investigate whether the joint bilateral filter could remove compression artefacts even when colour was both heavily subsampled and heavily compressed.

The Gaussian filter is known for its noise removal properties. Unfortunately, it blurs detail as well as noise. The bilateral filter<sup>9,10</sup> attempts to remedy this by introducing a further term that restricts ‘bleeding’ across image ‘edges’ by only blurring together pixels of similar colour or similar intensity:

$$v_p := \frac{\sum_{p' \in \Omega} g_d(p'-p) \times g_c(v_{p'} - v_p) \times v_{p'}}{\sum_{p' \in \Omega} g_d(p'-p) \times g_c(v_{p'} - v_p)} \quad (3)$$

where  $v_p$  is the value of pixel  $p$ ,  $p'-p$  is the Euclidean distance between pixels  $p'$  and  $p$ , and  $g_s$  is a Gaussian with zero mean and variance  $s$ .

This is essentially a Gaussian blur, bounded by ‘edges’ in the image, and normalised. This is useful when the noise is weaker than the edge information, but in heavily compressed chrominance layers we have little edge information, and in fact have spurious false edges introduced by JPEG’s lossy quantisation. However, we can blur the noisy image with respect to the luminance layer, using its high quality edge information. This is the joint bilateral filter<sup>5</sup>:

$$v_p := \frac{\sum_{p' \in \Omega} g_d(p'-p) \times g_c(R_{p'} - R_p) \times v_{p'}}{\sum_{p' \in \Omega} g_d(p'-p) \times g_c(R_{p'} - R_p)} \quad (4)$$


where  $R_p$  denotes the value of pixel  $p$  in the corresponding reference image, the luminance channel in our case. The joint bilateral is slow in operation in this form. A more efficient version is presented by Durand and Dorsey<sup>5</sup>, which uses a linear approximation and convolution. This is still too slow. Our improved version for the special case of chrominance upsampling can operate on a 1280×1024 image in under a minute, using a separability approximation. This performs a one-dimensional joint bilateral filter on each row of pixels, followed by a one-dimensional joint bilateral filter on each column of pixels. While the joint bilateral filter is not actually separable, this approximation produces results which are visually acceptable. This is partly because the human eye is so insensitive to the chrominance channels that the artefacts introduced by the approximation are generally good enough and partly because the other compression artefacts mask any artefacts which may be caused by the separation of the filter.

#### 5. USING THE JOINT BILATERAL FILTER

The standard JPEG algorithm encodes a colour image using the three channels Y, Cb and Cr, the first being luminance, the latter two chrominance. These are separately quantised (resulting in information loss) according to a ‘quality’ factor. In our new joint bilateral JPEG (JB-JPEG) algorithm we downsample chrominance more than in JPEG and upsample it on decompression using the joint bilateral filter with reference to the luminance channel. This allows us to reduce the size of the compressed chrominance channel, counterbalance this by improving the quality of the luminance channel, and hence achieve the same file size at standard JPEG with improved overall image quality.

In standard JPEG compression, the Cb and Cr (chrominance) layers can be downsampled, to exploit the fact that the eye is more sensitive to luminance and less sensitive to chrominance information. 4:1:1 sampling is usually the default setting, meaning four Y pixels (luminance) are stored for each Cb or Cr pixel. Each chrominance channel is thus scaled down by a factor of two in each of the two dimensions. This four to one downsampling is not generally noticeable. The joint bilateral filter, by contrast, facilitates downsampling of much higher factors.

Restoration of chrominance channels is performed by upsampling. The mechanism for this process is not standardized, but obvious methods are to use nearest neighbour, bilinear or bicubic interpolation. All these can give poor results when chrominance is downsampled by a factor higher than the standard default, producing either blocky artefacts (nearest neighbour) or colour bleeding across edges. However, the joint bilateral filter can give pleasing results with



**Fig. 11.** Left: original 512×512 images. Centre: chrominance downsampled to 26×26 (top) or 16×16 (bottom) then upsampled by bicubic interpolation. Note the blotches and colour bleeds. Right: chrominance downsampled by the same factor then upsampled by nearest-neighbour interpolation followed by joint bilateral filtering. Note the significant improvement in quality.

compression by a factor of over 400. This could offer a significant improvement to the JPEG compression mechanism. The results in Figure 11 highlight the relative insignificance of the chrominance layer.


The next stage of the JPEG algorithm is lossy quantisation. This introduces considerable artefacts, which render bicubic upsampling unsuitable as spurious patterns appear in the final image as colourful, and distracting, splodges. However, the joint bilateral filter is effective at removing these artefacts, as it tends only to preserve ‘real’ edges, by reference to the luminance channel’s edge detail, thus blurring away noise in the chrominance channels. Figure 12 gives an example.

## 6. RESULTS FOR JOINT BILATERAL FILTER

Thus far, we have shown that the joint bilateral filter allows us to reduce chrominance channel quality while retaining image appearance. We now make a direct comparison to JPEG to ascertain what sort of quality improvement can be achieved.

In our experiments, we generate two compressed files from the same source, both of approximately equal file size. One file is JPEG default with chrominance channels downsampled by a factor of four (2×2). The other is JB-JPEG with the chrominance channels downsampled by a factor of either 25 (5×5) or 100 (10×10). Chrominance is then compressed using JPEG at a high quality factor. However, the downsampling more than compensates for the storage requirements of high quality. As chrominance takes less file space in JB-JPEG, we increase the compression quality of the luminance channel to achieve roughly the same overall file size. This process is repeated for different quality settings (5, 10, 20, 40, and 50) of the standard JPEG algorithm. Figures 13–16 demonstrate example results, the types of improvement that this algorithm produces, and some of the remaining problems.

Our experiments show that the JB-JPEG approach always allows for some improvement in the luminance channel quality, but that this is at the cost of some loss of colour contrast. JB-JPEG improves luminance detail, reduces decompressed image noise, and can reduce size on disk. Its disadvantages are that it takes longer to process and that colour information is lost. In particular, some images exhibit noticeable loss of vivid colour information. Additional



**Fig. 12.** Typical JPEG quantisation at high compression introduces unappealing blocks of continuous colour, with visible noisy boundaries. The joint bilateral filter can remove this effect, smoothing out these regions whilst not losing edge detail. This makes it possible to use lossier quantisation, saving space further, without visibly sacrificing image quality.

processing, such as selective downsampling of low entropy regions, could offer further improvements to address these issues.


## 6. SUMMARY

Both methods correct for significantly higher subsampling on the chrominance channels than is attempted in the JPEG standard. These preliminary results indicate that such methods could be used to improve compressed image quality, while keeping file size constant, or decrease file size, while keeping quality constant. Their downside is their higher computational cost and the, as yet, poorly understood nature of the artefacts which they generate in decompressed imagery.

## REFERENCES

1. Pennebaker, W. B. & Mitchell, J. L., *JPEG: Still Image Data Compression Standard*, Springer, 1993.
2. Markle, W. & Hunt, B., "Coloring a black and white signal using motion detection. Canadian patent no. 1291260, Dec. 1987.
3. Welsh, T., Ashikhmin, M. & Mueller, K.. "Transferring colour to grayscale images", *ACM Trans. Graph.* **21**(3):277–280, 2002.






**Fig. 13.** Joint bilateral filter example. Left: JPEG compression using  $2 \times 2$  chrominance subsampling and IJG quality factor 20, file size 19,201 bytes. Right: JB-JPEG compression using  $5 \times 5$  chrominance subsampling, IJG quality factors of 22 for luminance and 75 for chrominance, file size 18,851 bytes. Joint bilateral filtering has prevented the colour shifts that are visible in the JPEG compressed image.


4. Levin, A., Lischinski, D. & Weiss, Y., “Colorization using optimization”, *ACM Trans. Graph.* **23**(3):689–694, 2004.
5. Durand, F. & Dorsey, J., “Fast bilateral filtering for the display of high-dynamic-range images”, *ACM Trans. Graph.* **21**(3):257–266, 2002.
6. Petschnigg, G., Szeliski, R., Agrawala, M., Cohen, M., Hoppe, H. & Toyama, K., “Digital photography with flash and no-flash image pairs”, *ACM Trans. Graph.* **23**(3):664–672, 2004.
7. Eisemann, E. and Durand, F., “Flash photography enhancement via intrinsic relighting”, *ACM Trans. Graph.* **23**(3):673–678, 2004.
8. Wallace, G. K., “The JPEG still picture compression standard”, *Communications of the ACM* **34**(4):30–44, 1991.
9. Tomasi, C. & Manduchi, R., “Bilateral filtering for grey and colour images”, *ICCV*, 836–846, 1998.
10. Smith, S. M. & Brady, J. M., “SUSAN—a new approach to low level image processing”, *IJCV* **23**:45–78, 1997.



**Fig. 14.** Joint bilateral filter example. Left: JPEG compression using  $2 \times 2$  chrominance subsampling and IJG quality factor 10, file size 8,053 bytes. Right: JB-JPEG compression using  $5 \times 5$  chrominance subsampling, IJG quality factors of 11 for luminance and 75 for chrominance, file size 7,822 bytes. As in Figure 13, joint bilateral filtering has prevented the colour shifts that are visible in the JPEG compressed image. The enlarged portion of the images shows that the JPEG compression has introduced noticeable and disturbing colour shift artefacts. Note, however, the large area of constant colour in the main image: this exhibits typically JPEG blocky artefacts in both images. These artefacts are present in the *luminance* channel and thus the JB-JPEG algorithm does not remove them.



**Fig. 15.** Joint bilateral filter example. Left: JPEG compression using  $2 \times 2$  chrominance subsampling and IJG quality factor 10, file size 16,762 bytes. Right: JB-JPEG compression using  $5 \times 5$  chrominance subsampling, IJG quality factors of 11 for luminance and 75 for chrominance, file size 16,590 bytes. Joint bilateral filtering has prevented the colour shifts that are visible in the JPEG compressed image. Note that the luminance channel artefacts that are easily visible in the magnified images are barely noticeable in the main images, while the chrominance artefacts in the JPEG image on left are noticeable in both the magnified and the main image.



**Fig. 16.** Joint bilateral filter example. Left: JPEG compression using  $2 \times 2$  chrominance subsampling and IJG quality factor 10, file size 10,453 bytes. Right: JB-JPEG compression using  $5 \times 5$  chrominance subsampling, IJG quality factors of 11 for luminance and 75 for chrominance, file size 10,345 bytes. Again notice that the joint bilateral filtering has prevented the colour shifts that are visible in the JPEG compressed image. However, note also that there are problems of colour bleeding in the JB-JPEG image where two adjacent areas of different *chrominance* have the same *luminance*. This is most obvious below the aircraft's red wing where the red is bleeding through onto the mountains.