

Summary A11, Research Skills

Title: “Image compression using sparse colour sampling combined with non-linear image processing”

The authors investigate if colour image compression could be addressed using two image-processing algorithms; colourisation and joint bilateral filter. Neither algorithm was designed for image compression. The aim is to find out if these algorithms could be used to improve compression rate for the same visual quality.

The first analysed algorithm is colourisation. This method is designed to colour a greyscale image using minimal amount of user intervention. A user gives colours only for small numbers of pixels. Then colours are propagated assuming that neighbourhood pixels with similar luminance probably have similar colour. In the compression by colourisation luminance is stored as standard, compressed, greyscale image and colour is stored at low sampling rate. Colour values are retained regularly spaced in the grid of greyscale pixels. Those colour and luminance values are used to reconstruct the image by colourisation algorithm.

The authors compare colourisation-based image compression against standard JPEG compression using PNSR values (PNSR = peak signal to noise ratio). In the experiments colourisation-based method generated fewer blocky artifacts, but the colours were more washed out. The experiments also showed that the degradation of image quality with respect to grid spacing is not uniform across all images. For many of the images colourisation-based method and JPEG had the same results. Colourisation worked better for than JPEG for the images with large smooth areas of

colour. JPEG was better for images with colour that varied artificially quickly.

Relationship between image quality and grid spacing may not be monotonic as observed for an image with large smooth monochromatic areas.

The second algorithm is joint bilateral filter. This method joins two original images with the same scene to produce better final image. In a modified JPEG decompression these two images are: luminance channel with high quality edge information and downsampled chrominance channels to match the edges in the luminance channel. In this new joint bilateral JPEG (JB-JPEG) algorithm chrominance is downsampled more than in JPEG. This allows improving the quality of the luminance for the same file size. Consequently, the image quality is improved. When chrominance is downsampled more than by a factor of 4, alternative ways to reconstruct it (nearest neighbourhood, bilinear, or bicubic interpolations) can give poor results, for example, blocky artefacts or colour bleeding across edges. Joint bilateral filter can overcome these problems even with compression by a factor of over 400. JPEG quantisation introduces artefacts, which appears in the final image as colourful splodges. Joint bilateral filter can remove these artifacts.

The authors compared JB-JPEG with JPEG. The experiment showed that overall JB-JPEG gives better visual results. It allows some improvement in the luminance channel at the expense of some loss of colour contrast. Additionally JB-JPEG reduces decompressed image noise, prevents colour shifts, and can reduce size on disk. The disadvantage is that the decompression of an image is slow. However, increases in computational power will make this method more feasible.

Summary B32, Research Skills

Title: “Balancing the expected and the surprising in geometric Op art”

The author takes a particular couple of 1960s works of art by the British artist Bridget Riley and analyses them to see whether he can find any meaningful patterns in them but most of the analysis of these two works of art are written up in a separate paper that gives the details of this analysis. The author then produces a range of alternative versions of the art work, using varying degrees of random perturbation, in a range of different algorithms with results that range in their attractiveness to the reader. The author produces three hypotheses about these types of art work.

The first hypothesis of the three is called “Distinguishing Types”. In this first hypothesis, the author claims that humans have considerable ability to distinguish between different types or different species or different algorithmic methods of generating patterns. The author claims that this hypothesis would be easy to prove using a method developed by some other researchers but then says that there is no point in trying to prove it as human beings can be trained to make arbitrary degrees of distinction, so that a beginner cannot tell the difference between any of the different species of patterns, but an expert can tell very very fine distinctions indeed.

The second hypothesis is called, by the author, “Aesthetic Balance”, by which the author means that the modification of the pattern away from a regular pattern towards an aesthetically interesting pattern can’t be purely random but must have

some sort of notion of aesthetic balance built into it. The author uses a range of prior work, principally by famous people (Alberti, Arnheim), to try to justify this position but, again, there is no proof of the concept and the author claims that further research is required to get anywhere with trying to work out how to encode this sort of idea in a computer program. He does not attempt to encode this in any way.

The third of the three hypotheses presented by the author is that of "Pattern Perception". The author, in this case, makes a firm prediction and undertakes a psycho-physical experiment to ascertain whether or not the prediction is correct. The prediction is that humans can easily detect patterns when less than 25% of the pattern is removed or disturbed, whereas removal of over fifty percent of the pattern destroys it. The author constructs an experiment with thirteen human subjects, four different patterns, and multiple runs of each pattern in order to see whether or not anything can actually be said about this hypothesis. The conclusion is that there is good evidence that humans can easily detect patterns if less than 25 percent has been removed and that humans have a very hard time indeed trying to detect patterns in more than 50% has been removed. The more interesting conclusion, that there is an aesthetically interesting region between these two values, is not really actually particularly amenable to mathematical analysis. The author asserts that we can very tentatively conclude that the evidence supports this rather interesting aesthetic conclusion.

Summary C42, Research Skills

Title: "A butterfly subdivision scheme for surface interpolation with tension control"

In the field of Computer Aided Geometric Design, the goal is to model smooth curves and surfaces. There exists two kinds of approaches. The basic approach is to model curves by giving a set of control points which form a polygon or a polyhedron, and a smoothing function, the second approach for modelling curves is quite different. The idea is to use the principal of binary subdivision. From a set of control points, the subdivision process generates recursively at each stage a new set containing twice as many elements. The process is called interpolatory if the new generated set contains all elements of the old set and in addition, new elements are inserted between old ones. Moreover, if the binary subdivision process is well defined, it converges to a smooth curve. For instance the four-point scheme is a well defined binary subdivision. This scheme is an interpolatory subdivision containing a tension parameter ω which defines tightening between resulting curve and initial polygone. This parameter offers a design flexibility and could be set to obtain different properties on the resulting curves. For instance, continuity, $C1$. Furthermore, by using the six-point scheme, we can obtain a $C2$ curves. There are already existing subdivision processes on surface but unfortunately, control points should be defined in regular square-like grid surface. The idea used to improve subdivision processes on surface is to generalise the four-point scheme. In every execution of the recursive process, each triangle face of the control polyhedron is transferred to four new triangle face by adding new points and refining

triangulation. New points are added with respect to the butterfly scheme which is a height-point rule, and new edges are added to refine triangulation. In addition the butterfly scheme has u, v, w as variables which can be set to produce a C^1 surface. Convergence analysis and surface regularity in subdivision scheme problem was already treated. Here are the results : a C^0 surface around a regular vertex of degree six could be obtained if the necessary condition $2u + 2v - 4w = 1$ holds. We can develop this condition (a) : $u = \frac{1}{2}, v = 2w$ which is the necessary condition to produce a C^1 surface around a regular vertex. Moreover, if there are no vertices of degree three in the triangulation and the condition $0 < \omega < \omega_0, \omega_0 > 1/16$ is satisfied, the resulting surface is C^1 . Finally a conjecture defines that if the condition a holds on irregular points and the surface around all other points is C^1 then the surface is globally C^1 . To get better flexibility design, the global variable ω is changed to a local variable which is associated to each point. Furthermore, a generalisation of the scalar tension parameter ω to a directed tension $\omega \in \mathbf{M}_{3 \times 3}$ allows a control of the tension parameter through all space directions.

A software has been developed to experiment the butterfly scheme. An experiment with a global tension parameter expresses a smooth surface and keeps points on three vertices. A second experimentation with local tension parameters $\omega = 0$ on a particular point T results to a stiff surface around T . Finally, we could observe fractal behaviour around T by doing the same experimentation with the directed tension parameter $\omega = \text{diag}(1/16, 1/16, 1/4)$ instead of the local tension parameter.

Summary D51, Research skills

Title: "A Time Sequential Multi-projector Autostereoscopic Display"

"A Time Sequential Multi-projector Autostereoscopic Display," written by Niel Dodgson et al, presents a three-dimensional display which is viewable by two people abreast at an optimal viewing distance. This display implements both stereo parallax and horizontal movement parallax, meaning that a different image is present to each eye and horizontal movement will result in new angles of observation.

The authors based the current design upon two previous methods for constructing 3-dimensional displays: the multi-projector display and the time-sequential display. The multi-projector display uses a set of independent projectors to project a series of images into a section of bands of space before the display such that each of the observer's eye sees a separate image, producing a stereo parallax, and movement of the head allows observation of a varying range of images, producing movement parallax.

The time-sequential display reduces cost by using a single-high-frame-rate display whose individual frames are multiplexed to be viewable in bands using a striped LCD mask and a Fresnel lens. The high-frame-rate monochromatic display can produce color by projecting three versions of each frame through a coloured LCD mask for every band of view. Thus, for the best colour time sequential display, which had 7 independent bands of view, twenty-one independent images must be displayed for each frame, requiring a refresh rate of 1200 hz for a 60HZ observed

frame rate. While the cost for this method is lower, the number of views is very limited.

The authors combine these two methods to produce a display capable of 28 bands of view. Four high-frame rate displays, each with its own striped LCD mask, are arranged behind a Fresnel lens. Each display projects 7 colour bands of view through the Fresnel lens. The four display regions fit together precisely such that 28 contiguous images are presented to the viewer. This system is less expensive than one utilising 28 independent projectors but has a much wider viewing angle than a single time-sequential display.

The resulting display allows ideal stereoscopic viewing 1,200 mm from the display over a breadth of about 2 feet. Each band of view is about 21 mm wide at this distance preventing both eyes from seeing the same view. Brightness and contrast were mediocre and the Fresnel lens picked up reflections from a wide angle, so the viewing area needed to be heavily hooded. Also the CRT phosphors were running at the upper limit of their refresh rate, so ghosting slightly degraded the stereo effect. Nevertheless, the display was usable and viewable from a broader angle than previous models.