Detection of colour filter array interpolated images

Andrew B. Lewis

UNIVERSITY OF CAMBRIDGE

Computer Laboratory

Digital imaging hardware

- Light from the subject passes through a lens and hits a charge-coupled device (CCD). The surface of the chip is an array of photoactive capacitors.
- After the exposure, a control circuit repeatedly causes each capacitor to transfer its charge to its neighbour, like a shift register. Each final capacitor transfers its charge to an amplifier which converts the charge into a voltage.
- An analog-to-digital converter samples the amplifier's output and stores it digitally in a bitmap.

Capturing a colour image

- To capture colour images, digital cameras often use a repetitive pattern of colour filters positioned over the capacitor array, called a colour filter array (CFA).
- For each square of four pixels, the CFA contains two green cells, one red cell and one blue cell.

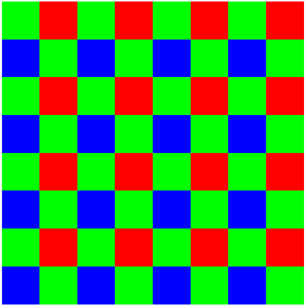

- Patented by Bryce E. Bayer of Eastman Kodak in 1976 (US patent 3,971,065).

Producing a red/green/blue image (demosaicing)

- The image capture device interpolates the two missing colour components at each pixel.
- Recall that $Y \approx 0.3R + 0.6G + 0.1B$. Green is often treated as the luma channel.

- *Bilinear* and *bicubic* interpolation apply to each channel independently. They convolve the input with a 2-D filter to find the missing values.
- *Smooth hue transition* interpolation applies bilinear interpolation to the green channel, then bilinearly interpolates the ratio R/G or B/G over missing red/blue pixels.
- *Median filter* interpolation calculates the bilinearly interpolated image, then applies a median filter to the pairwise differences between the channels ($R - G$, $R - B$, $G - B$).

CFA interpolation detection

Gallagher, Chen: Image authentication by detecting traces of demosaicing, Proc. CVPR WVU Workshop, 2008.

1. High-pass filter the green channel $\Rightarrow$ increase the difference in variance between original and interpolated samples.
2. The MLE of the variance of samples on a diagonal d is proportional to the mean of its absolute values[1], $m(d)$.
3. In interpolated images, this signal will be periodic over $T = 2$ samples $\Rightarrow$ peak in $\mathcal{F}_k\{m(d)\}_{d=0}^{k-1}$.
4. Use a threshold detector to check for a peak at this frequency (relative to the median value of the transformed signal).

[1] Assuming IID Gaussian samples

Terminology of inverse probability

Unknown parameters $\boldsymbol{\theta}$, data D, assumptions $\mathcal{H}$

$$P(\boldsymbol{\theta}|D,\mathcal{H}) = \frac{P(D|\boldsymbol{\theta},\mathcal{H})P(\boldsymbol{\theta}|\mathcal{H})}{P(D|\mathcal{H})}$$

$$\text{posterior} = \frac{\text{likelihood} \times \text{prior}}{\text{evidence}}$$

The quantity of $P(D|\boldsymbol{\theta},\mathcal{H})$ is a function of both D and $\boldsymbol{\theta}$. For fixed $\boldsymbol{\theta}$ it defines a probability over D. For fixed D it defines the likelihood of $\boldsymbol{\theta}$.

Maximum likelihood estimation

We wish to estimate $\boldsymbol{\theta}$ on the basis of data D. The maximum likelihood (ML) estimate of the parameters from the data is

$$\hat{\boldsymbol{\theta}}_{\mathsf{ML}}(D) = \arg\max_{\boldsymbol{\theta}} P(D|\boldsymbol{\theta}).$$

Pixel variance: bilinearly interpolation

$$\hat{X} = X * \frac{1}{4}\begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$

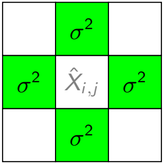

$$\mathrm{Var}(\hat{X}_{i,j}) = \begin{cases} \mathrm{Var}\left(\frac{1}{4}(\hat{X}_{i-1,j} + \hat{X}_{i,j-1} + \hat{X}_{i+1,j} + \hat{X}_{i,j+1})\right) = \frac{1}{4}\sigma^2 & \text{if } (i+j) \bmod 2 = 1, \\ \sigma^2 & \text{otherwise.} \end{cases}$$

Pixel variance: 2-D Laplace filtered (1)

$$Y = \hat{X} * \begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$

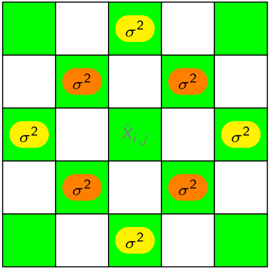

For non-interpolated pixels (i.e., $(i+j) \bmod 2 = 0$)

$$\begin{aligned}
\mathrm{Var}(Y_{i,j}) &= \mathrm{Var}\left(-4\hat{X}_{i,j} + \tfrac{1}{4}\left(\hat{X}_{i-2,j} + \hat{X}_{i-1,j+1} + \hat{X}_{i,j} + \hat{X}_{i-1,j-1} + \hat{X}_{i,j+2} + \hat{X}_{i+1,j+1} + \dots\right)\right) \\
&= \mathrm{Var}\left(-3\hat{X}_{i,j} + \tfrac{1}{2}\left(\hat{X}_{i-1,j-1} + \hat{X}_{i-1,j+1} + \hat{X}_{i+1,j+1} + \hat{X}_{i+1,j-1}\right)\right. \\
&\qquad \left. + \tfrac{1}{4}\left(\hat{X}_{i-2,j} + \hat{X}_{i,j+2} + \hat{X}_{i+2,j} + \hat{X}_{i,j-2}\right)\right) \\
&= 9\sigma^2 + \sigma^2 + \tfrac{1}{4}\sigma^2 = \tfrac{41}{4}\sigma^2, \\
\mathrm{E}(Y_{i,j}) &= -3\bar{X} + \tfrac{1}{2}\cdot 4\bar{X} + \tfrac{1}{4}\cdot 4\bar{X} = 0.
\end{aligned}$$

Pixel variance: 2-D Laplace filtered (2)

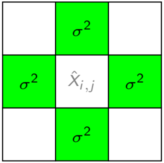

For interpolated pixels (i.e., $(i+j) \bmod 2 = 1$)

$$
\begin{aligned}
\text{Var}(Y_{i,j}) &= \text{Var}\left(-4 \cdot \tfrac{1}{4}(\hat{X}_{i-1,j} + \hat{X}_{i,j+1} + \hat{X}_{i+1,j} + \hat{X}_{i,j-1}) + \hat{X}_{i-1,j} + \hat{X}_{i,j+1} + \hat{X}_{i+1,j} + \hat{X}_{i,j-1}\right) \\
&= 0
\end{aligned}
$$

- The second-order differences of a bilinearly interpolated image have zero variance in interpolated pixels, and high variance in non-interpolated pixels.
- The algorithm works by estimating the variance along each diagonal (which consists entirely of interpolated or non-interpolated pixels). If the variances follow a periodic pattern down the image, this indicates that it may have undergone interpolation.

MATLAB source code

```matlab
function result = cfadetect(img)
    % Load the green channel of the image
    img = im2double(img) * 255.0;
    g = img(:, :, 2);

    % High pass filter
    laplace_matrix = [ 0  1  0; ...
                       1 -4  1; ...
                       0  1  0];
    filtered_g = conv2(g, laplace_matrix, 'valid');

    % Find the sum of all the diagonals in the filtered green channel
    diagonals = arrayfun(@(d) mean(abs(diag(filtered_g, d))),
                         -size(filtered_g, 1) + 1 : size(filtered_g, 2) - 1)

    % Show the DFT of the signal, with a log scale on the Y axis
    semilogy(abs(fft(diagonals)));

    result = filtered_g;
end
```