

DESIGNING COLOR FILTER ARRAYS FOR THE JOINT CAPTURE OF VISIBLE AND NEAR-INFRARED IMAGES

Yue M. Lu, Clément Fredembach, Martin Vetterli and Sabine Süsstrunk

School of Computer and Communication Sciences
Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland
Email: {yue.lu, clement.fredembach, martin.vetterli, sabine.sustrunk}@epfl.ch

ABSTRACT

Digital camera sensors are inherently sensitive to the near-infrared (NIR) part of the light spectrum. In this paper, we propose a general design for color filter arrays that allow the joint capture of visible/NIR images using a single sensor. We pose the CFA design as a novel spatial domain optimization problem, and provide an efficient iterative procedure that finds (locally) optimal solutions. Numerical experiments confirm the effectiveness of the proposed CFA design, which can simultaneously capture high quality visible and NIR image pairs.

Index Terms— Color filter arrays, color acquisition, sampling, near-infrared, digital photography

1. INTRODUCTION

Silicon-based camera sensors are inherently sensitive to the near-infrared (NIR) band of the light spectrum. Most camera designs place an NIR-blocking filter, usually named “hot mirror”, in front of the sensor to prevent the NIR contamination of the visible image. By filtering out such a large part of the spectrum (700 nm–1100 nm), however, a significant amount of potentially valuable information is lost.

The usefulness of the NIR band has long been recognized in fields such as spectroscopy [1] and remote sensing [2]. Providing complementary information to the visible data, the NIR band has also shown great potential in vision and digital photography applications. For example, recent work [3]–[6] has demonstrated that the joint processing of visible and NIR data results in image enhancement and analysis capabilities beyond what can be achieved using visible information only.

The major limitation of research and applications of joint visible/NIR image processing is the current acquisition method: One simple way is to capture the visible and NIR images of the same scene sequentially, a manner that is both cumbersome and artifact-prone (due to camera and/or scene movement between consecutive shots). An alternative is to use a two-camera rig together with a beam-splitter [3], an accurate but fairly expensive setup suitable only for certain professional applications.

Motivated by the above issue, we are currently investigating a camera design that can *simultaneously* capture high-quality visible/NIR image pairs with a *single* sensor. The focus of the present paper is to address one critical component of such a camera—the color filter array (CFA)—in detail. We pose

the CFA design as a spatial-domain optimization problem, and present an efficient iterative procedure that finds (locally) optimal solutions. Numerical experiments demonstrate the feasibility and effectiveness of the proposed design.

2. BACKGROUND

2.1. Color Filter Arrays

A camera sensor is, in essence, monochromatic. In order to acquire color information, the preferred solution is to place a CFA in front of the sensor. This array is composed of a mosaic of colored transparent material that allows only a portion of the spectrum to pass through.

In general, a CFA can be represented by a triplet

$$\mathbf{cfa}[\mathbf{n}] = [c_r[\mathbf{n}], c_g[\mathbf{n}], c_b[\mathbf{n}]]^T \in [0, 1]^3,$$

where the three components denote the relative percentages of the R, G, and B information retained at pixel location \mathbf{n} . Note that we enforce the range of CFA values to be within $[0, 1]$ to ensure physical realizability (via subtractive color layers). Let $r[\mathbf{n}]$, $g[\mathbf{n}]$ and $b[\mathbf{n}]$ denote the ground truth R, G, B values of the scene (i.e., the ones we would obtain from a 3-CCD camera); the sensor reading after the CFA can then be modeled as

$$y_{\text{vis}}[\mathbf{n}] = c_r[\mathbf{n}]r[\mathbf{n}] + c_g[\mathbf{n}]g[\mathbf{n}] + c_b[\mathbf{n}]b[\mathbf{n}]. \quad (1)$$

2.2. Near-Infrared Acquisition

We propose a camera design for the simultaneous capture of high quality visible/NIR images, which requires no modification of the current imaging sensors. In fact, standard silicon-based sensors, both CCD and CMOS, are intrinsically sensitive to wavelengths from roughly 200 nm to 1100 nm. Thus, if one removes the hot mirror from the camera, the sensors will have the capability of imaging both the visible and NIR bands.

One component of the camera that does need modification is the CFA. Without the hot mirror, the CFA filters are partly transparent not only to their respective color wavelengths, but to NIR as well [4]. This fact makes the acquisition model in (1) no longer applicable.




Fig. 1. The block diagram for the visible/NIR acquisition process in the proposed camera design.




Fig. 2. A generic example of a 2×2 periodic CFA.

Assume that all color filters in the CFA are equally transmissive to the NIR spectrum.¹ Consequently, the color/NIR acquisition process in the proposed camera design can be represented by the scheme shown in Figure 1. The three visible channels, $r[\mathbf{n}]$, $g[\mathbf{n}]$ and $b[\mathbf{n}]$, are combined by the CFA into a single mosaicked image $y_{\text{vis}}[\mathbf{n}]$. However, since the hot mirror of the camera is removed, all color filters are sensitive to the NIR spectrum as well. The actual sensor reading $y[\mathbf{n}]$ therefore is a summation of $y_{\text{vis}}[\mathbf{n}]$ and an NIR intensity image, denoted by $f[\mathbf{n}]$. It follows from (1) that

$$y[\mathbf{n}] = c_r[\mathbf{n}]r[\mathbf{n}] + c_g[\mathbf{n}]g[\mathbf{n}] + c_b[\mathbf{n}]b[\mathbf{n}] + f[\mathbf{n}]. \quad (2)$$

The goal of the present work is to design a CFA and the associated demosaicking algorithm that can reconstruct, from $y[\mathbf{n}]$, full-resolution estimates of the visible and NIR information (i.e., a 4-layer RGB+NIR image).

Before presenting the proposed new design in Section 3, we first explain why existing CFA patterns in the literature cannot handle the simultaneous acquisition of the visible/NIR images. To that end, consider a simple image

$$r[\mathbf{n}] = g[\mathbf{n}] = b[\mathbf{n}] = c_1 \text{ and } f[\mathbf{n}] = c_2,$$

where c_1 and c_2 are two constants. In the visible spectrum, the above image represents a uniform gray patch. A common feature of most existing CFA patterns is that

$$c_r[\mathbf{n}] + c_g[\mathbf{n}] + c_b[\mathbf{n}] \equiv \gamma, \quad (3)$$

i.e., the summation of the R, G, B values (sometimes called the “luminance gain”) of the CFA is a constant. For example, the widely used Bayer CFA satisfies this condition with $\gamma = 1$. Other examples include the CMY CFA ($\gamma = 2$), as well as the various new designs proposed in [7]. From (2), the sensor reading in this case can be written as $y[\mathbf{n}] = \gamma c_1 + c_2$. We see that there exist different combinations of c_1 and c_2 that can generate identical $y[\mathbf{n}]$. Consequently, it is impossible to recover the original images by using CFAs satisfying (3).

¹Pigments with this property do exist, and are used, for example, in the printing of banknotes.

3. THE PROPOSED CFA DESIGN METHOD

The choice of CFA patterns has a great impact on the final image quality in the digital imaging pipeline. The recent work of Hiraoka and Wolfe [7] proposes to design the CFA in the Fourier domain. The key idea behind this Fourier approach is that one should design the CFA to minimize the frequency-domain aliasing between the luminance and chrominance channels [8].

By doing so, aliasing is treated as noise, and hence should be avoided at all cost. However, rather than a total loss of information, aliasing merely represents a *linear* mixing of frequency values, which can be subsequently decoupled by carefully designed reconstruction algorithms. In what follows, we propose a novel spatial domain approach to CFA design, which allows the existence of frequency aliasing.

3.1. Linear Minimum Mean Square Error Demosaicking

We start our discussion with the demosaicking algorithm. Consider a generic $N \times N$ periodic CFA. Surrounding each $N \times N$ block is a local neighborhood of pixels of size $(2L + 1)N \times (2L + 1)N$, where L specifies the neighborhood size. See Figure 2 for an example, where $N = 2$ and $L = 1$.

Denote by $\mathbf{y} \stackrel{\text{def}}{=} \text{vec}(y[\mathbf{n}])$ the sensor observation vector, where $\text{vec}(\cdot)$ denotes the vectorization of a matrix by stacking its columns. Similarly, we can define the ground truth vector $\mathbf{x} \stackrel{\text{def}}{=} [\text{vec}(r[\mathbf{n}])^T, \text{vec}(g[\mathbf{n}])^T, \text{vec}(b[\mathbf{n}])^T, \text{vec}(f[\mathbf{n}])^T]^T$. The proposed visible/NIR image acquisition model in (2) can then be written as a compact matrix-vector multiplication

$$\mathbf{y} = \mathbf{A} \mathbf{x} \stackrel{\text{def}}{=} [\text{diag}(c_r[\mathbf{n}]) \quad \text{diag}(c_g[\mathbf{n}]) \quad \text{diag}(c_b[\mathbf{n}]) \quad \mathbf{I}] \mathbf{x}, \quad (4)$$

where $\text{diag}(\cdot)$ denotes a diagonal matrix constructed from its argument, and \mathbf{I} is an identity matrix. We refer to the matrix \mathbf{A} defined above as the *CFA sampling matrix*.

Let $\mathbf{x}_0 \in \mathbb{R}^{4N^2}$ be the vector formed by the RGB and NIR pixel values at the center $N \times N$ block (e.g., the region inside the black rectangle in Figure 2). We can easily verify that there is a constant “selection” matrix \mathbf{S} (consisting of zeroes and ones) such that $\mathbf{x}_0 = \mathbf{S} \mathbf{x}$.

The goal of demosaicking is to obtain an estimate of \mathbf{x}_0 from the observation vector \mathbf{y} . In this paper, we focus on linear demosaicking algorithms, and hence the estimation process can be represented by

$$\hat{\mathbf{x}}_0 = \mathbf{D} \mathbf{y} = \mathbf{D} \mathbf{A} \mathbf{x}, \quad (5)$$

where $\hat{\mathbf{x}}_0$ is the estimated vector and \mathbf{D} is a fixed demosaicking matrix. The average performance of a particular demosaicking matrix \mathbf{D} can be measured by the mean square error of reconstruction, defined as

$$\text{MSE} \stackrel{\text{def}}{=} \text{E}(\|\mathbf{x}_0 - \hat{\mathbf{x}}_0\|^2) = \text{E}(\|\mathbf{S} \mathbf{x} - \mathbf{D} \mathbf{A} \mathbf{x}\|^2), \quad (6)$$

where $\text{E}(\cdot)$ denotes the expectation operator.

Let $\mathbf{C} \stackrel{\text{def}}{=} \text{E}(\mathbf{x} \mathbf{x}^T)$ be the data correlation matrix. Since \mathbf{C} is positive semidefinite, we can always factorize it as $\mathbf{C} =$

PP , where the “square root” P is another positive semidefinite matrix. The MSE defined in (6) can be rewritten as [9]

$$\text{MSE} = \|SP - DAP\|_F^2, \quad (7)$$

where $\|\cdot\|_F$ is the Frobenius norm of a matrix.

The optimal linear demosaicking scheme in the minimum mean square error (MMSE) sense is thus the solution to the following optimization problem: [10]

$$D^* = \arg \min_D \|SP - DAP\|_F^2. \quad (8)$$

A closed-form solution to (8) is

$$D^* = SP(AP)^\dagger, \quad (9)$$

where \dagger denotes the pseudo-inverse of a matrix.

3.2. A Spatial Domain Approach to Optimal CFA Design

For an $N \times N$ periodic CFA, its color pattern within one period is specified by $3N^2$ numbers

$$\{c_r[\mathbf{n}], c_g[\mathbf{n}], c_b[\mathbf{n}] : \mathbf{n} \in [0, N-1]^2\}. \quad (10)$$

We denote by $\alpha \in \mathbb{R}^{3N^2}$ the column vector containing all $3N^2$ color values in (10). As stated before, we only consider $\alpha \in [0, 1]^{3N^2} \stackrel{\text{def}}{=} \mathcal{B}$ to ensure physical realizability. Note that the CFA sampling matrix defined in (4) can now be written as $A(\alpha)$, i.e., a (matrix-valued) function of α .

Given the LMMSE demosaicking scheme in (8), we propose that the optimal CFA pattern is the solution to the following double optimization problem

$$\alpha^* = \arg \min_{\alpha \in \mathcal{B}} \left(\min_D \|SP - DA(\alpha)P\|_F^2 \right). \quad (11)$$

There exist close analogies between the above task and several classical problems in communication and learning theories. In particular, it is helpful to interpret the CFA sampling matrix A in (11) as a low-dimensional approximation operation and, correspondingly, demosaicking D as the best reconstruction operation. Finding the optimal CFA is thus equivalent to finding the optimal approximation scheme for the original signals with minimum information loss.

A closed-form solution of (11) does not seem to exist in general. Instead, we employ and extend the alternating minimization algorithm proposed in [9], which can find (locally) optimal solutions. To start, we first rewrite the “sequential” optimization problem in (11) into the following “simultaneous” optimization scheme

$$(\alpha^*, D^*) = \arg \min_{\alpha \in \mathcal{B}, D} \|SP - DA(\alpha)P\|_F^2. \quad (12)$$

See [11] for a justification of the equivalence of (11) and (12).

For fixed α , the above problem is convex, and the corresponding optimal solution $D^*(\alpha)$ is given by (9). Now consider

fixing D and searching for the best α . The key observation is that we can rewrite $A(\alpha)$ as a linear combination

$$A(\alpha) = A_0 + \sum_{k=1}^{3N^2} \alpha_k A_k, \quad (13)$$

where $\{A_k\}_{k=0,1,\dots,3N^2}$ are constant matrices whose entries can be determined by (4). Substituting (13) into (12), we get

$$\alpha^*(D) = \arg \min_{\alpha \in \mathcal{B}} \|SP - DA_0P - \sum_{k=1}^{3N^2} (\alpha_k DA_kP)\|_F^2. \quad (14)$$

This is a quadratic programming problem with inequality constraints (since $\alpha \in \mathcal{B}$). It can be efficiently solved by methods such as the interior point algorithm. We can now summarize the proposed iterative search procedure as follows.

Procedure 1 (Iterative Search for Optimal CFAs) *Start from an initial guess of the CFA (i.e., a vector $\alpha \in \mathcal{B}$).*

1. *Initialize: $i = 1$ and $e^{(0)} = 0$.*
2. *Compute $D = SP(A(\alpha)P)^\dagger$.*
3. *Calculate the MSE $e^{(i)} = \|SP - DA(\alpha)P\|_F^2$.*
4. *Solve the constrained quadratic minimization problem (14) and set the solution to α .*
5. *If $|e^{(i)} - e^{(i-1)}|$ is greater than a given threshold (e.g., $\delta = 0.001$), then set $i \leftarrow i + 1$ and return to Step 2. Otherwise, stop the procedure and return α .*

4. NUMERICAL EXPERIMENTS

We present numerical experiments in this section to verify the effectiveness of the proposed CFA design. To obtain the test images, we modified a Canon 300D camera by replacing its hot mirror with a piece of clear glass. This modification allows the camera to capture visible and NIR light at the same time. We use lens-mounted filters to capture the visible and NIR images of the same scene in two consecutive shots. Image registration has been applied to the visible/NIR pair to compensate for the relative camera movement between the two shots. To further reduce the artifacts due to remaining registration errors and other in-camera processing, we downsample the original images from 2000×3000 pixels to 512×768 pixels.

In our experiments, we use a total of 12 visible/NIR image pairs, of which six are used as the training set for estimating the data correlation matrix C . In what follows, we present the results for one set of parameters: $N = 4$ (i.e., 4×4 CFA patterns) and $L = 1$ (i.e., a neighborhood size of 12×12). The selection process of these parameters is omitted due to space limitations.

Figure 3(a) shows the convergence behavior of the alternating minimization algorithm proposed in Procedure 1, with a randomly generated starting point. The MSE values decrease monotonically throughout the iteration process, and the entire procedure converges within a small tolerance ($\Delta\text{MSE} < 0.001$)




Fig. 3. (a) The convergence of the proposed iterative procedure. (b) The obtained 4×4 CFA design with 15 visible/NIR filters and one NIR-only filter (i.e., the black pixel on the first row).

after about 1650 iterations. To improve the chance of reaching the global minimum, we repeat the iterative procedure 20 times, with different randomly generated initial values. The obtained 4×4 CFA with the smallest MSE is shown in Figure 3(b).

The average reconstruction MSE for the six test image pairs (not the training set) are 17.4 and 16.4, for visible and NIR, respectively. Note that if we only acquire the visible part, the same set of images will lead to an average MSE of 12.1 (using the visible-only CFA in [9]). Effectively, we trade spatial resolution for additional spectral information. Considering the potential applications of capturing the NIR alongside the visible images, and the fact that current sensor resolutions have increased beyond the human visual system’s discriminating capabilities, this is a trade we can easily afford.

Figure 4 displays the comparison between two pairs of original RGB/NIR images and the demosaicked results. We observe that, in both the visible and the NIR channels, all key visual information of the original images have been faithfully retained, demonstrating the feasibility of acquiring high quality visible and NIR images simultaneously using a single sensor.

5. CONCLUSIONS

We proposed in this paper a general method for designing CFAs that can simultaneously capture visible and NIR images. Unlike previous work on CFA design which approaches the problem in the Fourier domain, we pose the CFA design as a novel spatial domain optimization problem. We provide an efficient iterative procedure to search for CFAs that are (locally) optimal solutions. Numerical experiments confirm the effectiveness of the proposed CFA design, which can faithfully retain the information of the original visible and NIR images.

6. REFERENCES

- [1] D. A. Burns and E. W. Ciurczak, *Handbook of Near-Infrared Analysis*, CRC Press, 2nd edition, 2001.
- [2] T. M. Lillesand and R. W. Kiefer, *Remote Sensing and Image Interpretation*, Wiley, 4th edition, 1999.
- [3] X.-P. Zhang, T. Sim, and X.-P. Miao, “Enhancing photographs with near infrared images,” in *Proc. IEEE Int. Conf. on Computer Vision and Pattern Recognition*, Ankorage, US, 2008.




Fig. 4. Left column: two pairs of original RGB/NIR images. Right column: the corresponding demosaicked images using the proposed CFA and the linear demosaicking scheme in (8).

- [4] C. Fredembach and S. Ssstrunk, “Colouring the near-infrared,” in *Proc. IS&T/SID 16th Color Imaging Conference*, 2008, pp. 176–182.
- [5] C. Fredembach and S. Ssstrunk, “Illumination estimation and detection using near-infrared,” in *Proc. SPIE Electronic Imaging, Digital Photography V*, 2009.
- [6] L. Schaul, C. Fredembach, and S. Ssstrunk, “Color image de-hazing using the near-infrared,” in *Proc. IEEE Int. Conf. on Image Proc.*, Cairo, Egypt, 2009.
- [7] K. Hirakawa and P. J. Wolfe, “Spatio-Spectral color filter array design for optimal image recovery,” *IEEE Trans. Image Process.*, vol. 17, no. 10, pp. 1876–1890, Oct. 2008.
- [8] D. Alleysson, S. Ssstrunk, and J. Hraut, “Linear demosaicing inspired by the human visual system,” *IEEE Trans. Image Process.*, vol. 14, no. 4, pp. 439–449, Apr. 2005.
- [9] Y. M. Lu and M. Vetterli, “Optimal color filter array design: Quantitative conditions and an efficient search procedure,” in *Proc. SPIE Electronic Imaging, Digital Photography V*, 2009.
- [10] B. Chaix de Lavarne, D. Alleysson, and J. Hraut, “Practical implementation of LMMSE demosaicing using luminance and chrominance spaces,” *Computer Vision and Image Understanding*, vol. 107, no. 1-2, pp. 3–13, 2007.
- [11] G. H. Golub and V. Pereyra, “The differentiation of pseudo-inverses and nonlinear least squares problems whose variables separate,” *SIAM J. Numer. Anal.*, vol. 10, no. 2, pp. 413–432, Apr. 1973.