Mahesh B et al., Vol 3(1), pp 86-100
https://doi.org/10.55306/CJDTES.2026.030202
e-ISSN: 3048-8389
Ci-STEM Journal of Digital Technologies and Expert Systems

From Classical Interpolation to Deep Neural Reconstruction: A Comprehensive Survey of Colour Filter Array Demosaicking Algorithms

Raja Rao Chatla[1], Soumitra Kumar Mandal[2]
[1]Department of Electrical Engineering, Maulana Abul Kalam Azad University of Technology, Kolkata, West Bengal, India
[2]Department of Electrical Engineering, National Institute of Technical Teachers' Training & Research (NITTTR), Kolkata, West Bengal, India.
[1]email: crrao@bopter.gov.in, [2]email: skmandal@nitttrkol.ac.in.

Open Access
Review Article
Received: 22/06/2026
Accepted: 05/08/2026
Published: 27/08/2026
Corresponding author email:
crrao@bopter.gov.in
Citation:
C. R. Rao et al., "From Classical Interpolation to Deep Neural Reconstruction: A Comprehensive Survey of Colour Filter Array Demosaicking Algorithms" Ci-STEM Journal of Digital Technologies and Expert Systems, Vol. 3(2), pp. 86-100, 2026
doi: 10.55306/CJDTES.2026.030202
Copyright:
©2026 Mahesh B et al.,
This is an open-access article distributed under the terms of the Creative Commons Attribution License which grants the right to use, distribute, and reproduce the material in any medium, provided that proper attribution is given to the original author and source, in accordance with the terms outlined by the license.
(https://creativecommons.org/licenses/by/4.0/)
Published By:
Ci-STEM Global Services Foundation, India.

Abstract:
Colour filter array (CFA) demosaicking is the process by which a single-sensor digital camera reconstructs a full-resolution, three-channel colour image from a spatially sub-sampled mosaic of red, green and blue measurements. Because nearly every consumer camera, smartphone and machine-vision sensor in use today relies on some variant of a Bayer, Quad-Bayer or Nona-Bayer colour filter array, the quality of the demosaicking step places a hard ceiling on the achievable image quality of the entire imaging pipeline. Early work on this problem, spanning roughly 1987 to the mid-2000s, produced a rich family of model-based interpolation schemes: nearest-neighbour replication, bilinear and smooth-hue-transition interpolation, median-based and edge-sensing algorithms, block-matching, edge-directed interpolation, and directional-filtering-with-a-posteriori-decision schemes. These methods remain instructive because they encode, in closed analytic form, the core priors that all later data-driven methods must eventually learn: spectral correlation between colour channels, local edge orientation, and piecewise smoothness. Since approximately 2016, the field has been reshaped by deep learning. Convolutional networks, residual and iterative denoising cascades, attention and Transformer-based architectures, implicit neural representations, and, most recently, state-space (Mamba) and binarized hybrid networks have driven reported PSNR figures on standard benchmarks up by eight to fourteen decibels relative to bilinear interpolation, while simultaneously extending the task to jointly address denoising, non-Bayer sensor layouts (Quad-Bayer, Nona-Bayer, HybridEVS), and real-time mobile deployment constraints. This paper presents a comprehensive survey that (i) revisits the classical algorithmic foundations of demosaicking, (ii) systematically reviews the deep-learning literature from 2016 through 2026, organised by architectural paradigm, (iii) consolidates benchmark datasets and evaluation protocols used across this literature, (iv) compiles and compares quantitative results reported for representative algorithms on the Kodak and McMaster datasets, and (v) discusses open challenges — including generalisation across sensor patterns, joint demosaicking-denoising under real noise, and efficient on-device inference — together with promising future research directions such as unified multi-CFA models, diffusion-based priors, and self-supervised training regimes..
Keywords:
Colour Filter Array, Demosaicking, Bayer Pattern, Quad-Bayer, Nona-Bayer, Convolutional Neural Networks, Vision Transformers, State-Space Models, Image Signal Processing, Joint Denoising-Demosaicking.

1. INTRODUCTION

A digital colour photograph is, at the point of capture, not actually a colour image at all. Almost every consumer camera, smartphone, industrial camera and machine-vision sensor manufactured today uses a single monochrome CCD or CMOS sensor overlaid with a mosaic of miniature colour filters, so that any individual photosite records the intensity of only one of the three primary colours — red, green or blue. Reconstructing a full-resolution RGB image from this sparsely and irregularly sampled mosaic is the job of the demosaicking (also spelled "demosaicing" or "debayering") algorithm, and it is one of the first, and most consequential, steps in the modern camera image signal processing (ISP) pipeline. Every subsequent stage — white balance, tone mapping, sharpening, compression, and any downstream computer-vision task — inherits the errors introduced at this stage, so the accuracy of demosaicking has a direct and measurable effect on final image quality.

The engineering motivation for single-sensor colour capture is straightforward: recording red, green and blue values at every spatial location would require either three co-registered sensors and a beam-splitting prism assembly (used in some broadcast cameras but far too bulky, heavy and expensive for

consumer devices) or three vertically stacked photodiodes per pixel (used in a small number of specialised sensors such as Foveon X3, but difficult to scale economically). The dominant commercial solution, patented by Bryce Bayer at Kodak in 1976, instead places a single colour filter over each photosite in a repeating mosaic, and estimates the two missing colour values at every pixel computationally.

1.1 The Colour Filter Array and the Bayer Pattern

In the classical Bayer colour filter array, green filters are arranged on a quincunx (diagonal) lattice occupying half of all photosites, while red and blue filters occupy one quarter each of the remaining locations. The disproportionate allocation to green is deliberate: the luminance response of the human visual system peaks in the same part of the spectrum as the green channel, and the eye is considerably more sensitive to spatial detail in luminance than in chrominance. By sampling green at twice the density of red and blue, the sensor allocates more of its limited spatial-sampling budget to the channel that carries the most perceptually important high-frequency detail, while red and blue — which behave more like chrominance information — can be sub-sampled more aggressively without a proportional loss in perceived sharpness.

Figure 1 illustrates the Bayer layout together with two increasingly common non-Bayer alternatives. Demosaicking is fundamentally an interpolation and inverse-problem task: given the single-channel mosaic, the algorithm must estimate, at every pixel, the two colour channels that the sensor did not measure at that location, while avoiding two characteristic classes of artefact. The first is the "zipper effect", a jagged, false alternation of colour along high-contrast edges that results from naive linear interpolation smoothing across an edge in one channel but not another. The second is false colour or colour moiré, a low-frequency chromatic aliasing artefact that appears in regions of fine, high-frequency texture (such as fabric weaves or distant foliage) where the mosaic sampling rate is insufficient relative to the scene's spatial frequency content.

1.2 Beyond Bayer: Quad-Bayer, Nona-Bayer and HybridEVS Sensors

Smartphone image sensors have, since roughly 2019, increasingly departed from the single-Bayer layout in pursuit of improved low-light performance. Quad-Bayer sensors group each colour filter into a 2×2 block of identically coloured photosites, allowing four adjacent photosites of the same colour to be binned into a single, larger, lower-noise "super-pixel" in low light, while the same sensor can be read out at full native resolution and re-mosaicked to an effective Bayer pattern in bright light. Nona-Bayer sensors extend the same idea to 3×3 blocks. More recently, hybrid sensors that interleave conventional RGB photosites with dedicated event-sensing pixels — referred to in the literature as HybridEVS or Quad-Bayer HybridEVS sensors — combine frame-based colour imaging with asynchronous event-camera-style motion sensing on a single silicon die, and introduce an even more irregular, sensor-specific mosaic pattern that conventional demosaicking algorithms were never designed to handle. This proliferation of sensor layouts is one of the central forces reshaping demosaicking research in the 2020s, and is discussed at length in Section 3.7.

1.3 Motivation and Scope of this Survey

The demosaicking literature now spans nearly four decades and several distinct methodological eras: heuristic pixel-level interpolation (late 1980s to late 1990s), edge-aware and frequency-domain methods (late 1990s to 2010), sparse and model-based regularisation methods (2010 to 2016), and the deep-learning era (2016 to the present), within which convolutional, attention-based, implicit-neural and state-space architectures have each, in turn, defined the state of the art. This survey has two goals. First, it consolidates the classical algorithmic foundations — several of which remain the reference baselines against which every learned method is still benchmarked — into a single coherent narrative. Second, and more centrally, it provides a systematic, comparative review of demosaicking research published from 2016 through mid-2026, organised by architectural paradigm, together with a consolidated, source-cited comparison of quantitative results reported across this body of work. Unlike narrower prior surveys that focus on a single family of methods, this paper explicitly bridges the classical and learning-based literatures and extends the discussion to the non-Bayer, joint denoising-demosaicking, and resource-constrained deployment settings that dominate current publication activity.

The remainder of the paper is organised as follows. Section 2 reviews the principal classical, model-based demosaicking algorithms. Section 3 surveys the deep-learning literature across nine architectural

sub-paradigms. Section 4 catalogues the benchmark datasets used throughout this literature, and Section 5 reviews the associated evaluation metrics. Section 6 presents a consolidated, source-cited quantitative comparison across representative algorithms. Section 7 discusses cross-cutting challenges, Section 8 outlines promising future research directions, and Section 9 concludes the paper.

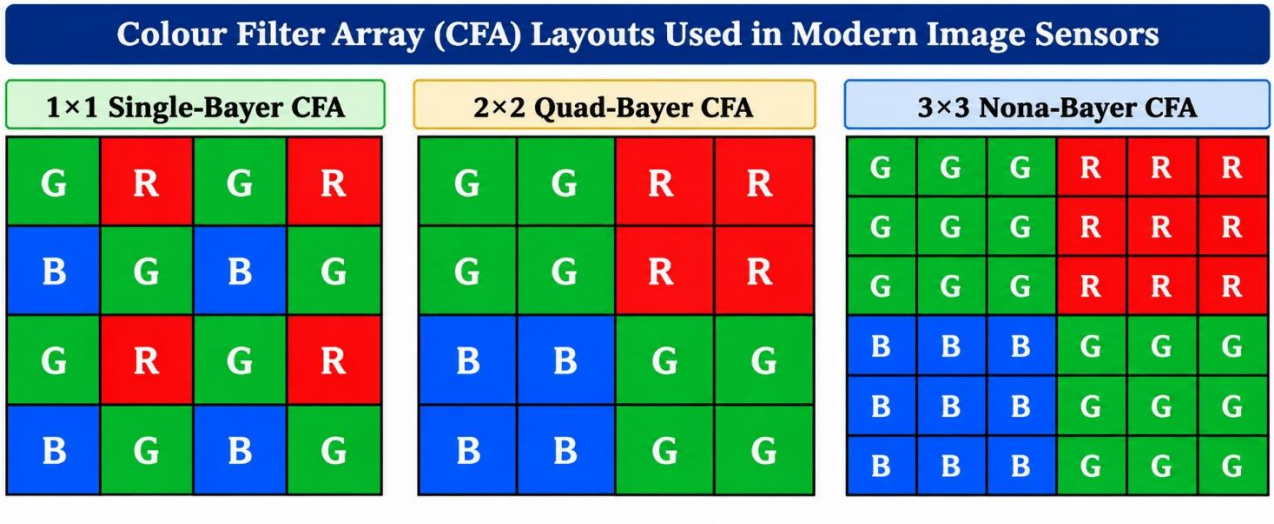

Figure 1. Common colour filter array (CFA) layouts: 1×1 Single-Bayer, 2×2 Quad-Bayer, and 3×3 Nona-Bayer, illustrating the increasing spatial grouping of same-colour photosites used for pixel binning in low-light conditions.

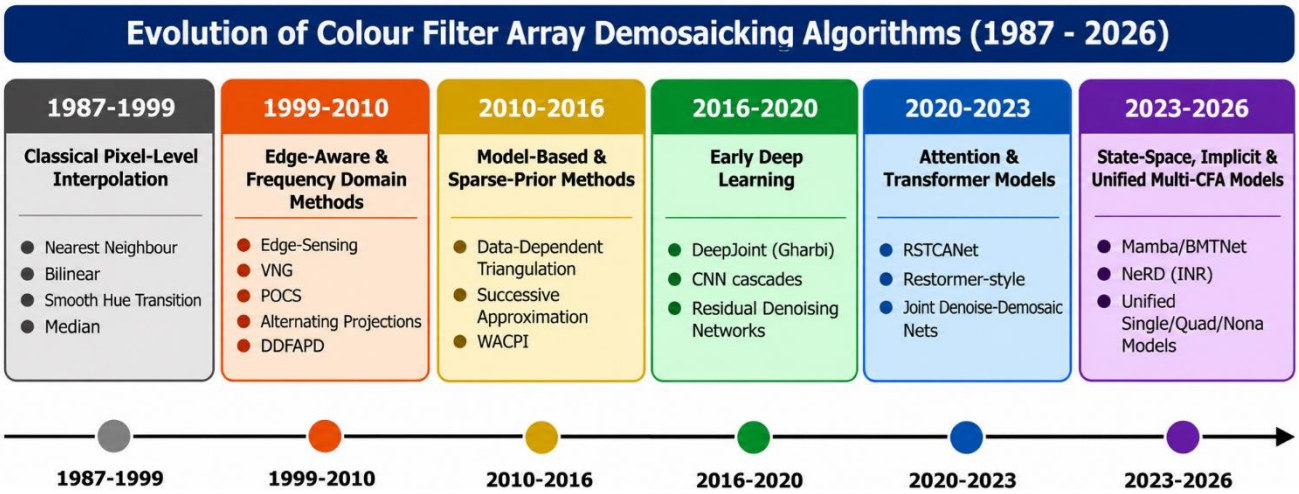

Figure 2. Approximate methodological eras in demosaicking research, from early heuristic pixel-level interpolation through the current state-space / implicit-representation era. Boundaries are approximate and overlap in practice.

2. CLASSICAL, MODEL-BASED DEMOSAICKING ALGORITHMS

Model-based demosaicking algorithms estimate missing colour samples using hand-designed rules that encode assumptions about local image structure — typically some combination of spectral correlation between colour channels and local smoothness or edge orientation. Although these methods have been superseded in raw reconstruction accuracy by learned models, they remain important for three reasons: they require no training data, they run with predictable, low computational and memory cost suitable for in-sensor or embedded hardware, and they continue to serve as the standard baselines and, in several learned pipelines, as an initialisation step or hybrid post-processing stage. Table 1 summarises the principal classical algorithms in roughly chronological order of introduction.

2.1 Pixel-Level and Bilinear Methods

The simplest approach, nearest-neighbour replication (Ozawa, 1987), assigns each missing colour sample the value of an adjacent measured pixel of the same channel, without any interpolation weighting. Bilinear interpolation (Hamilton and Adams, 1997) improves on this by averaging the two

or four nearest same-colour neighbours, but because it interpolates each colour plane independently, it is prone to pronounced zippering at edges and desaturation in high-frequency texture. Smooth hue transition interpolation (Cok, 1987), available in both linear and logarithmic exposure-space variants, was among the first schemes deployed in commercial cameras; it interpolates the luminance-like green channel first using bilinear interpolation, then reconstructs red and blue by imposing a smooth "hue" ratio or difference relative to the already-interpolated green channel, which substantially reduces the abrupt hue shifts characteristic of naive channel-independent bilinear interpolation.

2.2 Edge-Sensing and Directional Methods

Because uniform averaging across an edge is the principal source of zippering artefacts, a second generation of algorithms introduced explicit edge detection into the interpolation decision. Median-based interpolation (Freeman, 1988) performs an initial bilinear pass and then median-filters the colour-difference planes (R−G and B−G), which preserves edges considerably better than bilinear interpolation alone while remaining computationally light. Edge-sensing interpolation algorithms (Laroche and Prescott, 1994; Hibbard, 1995) compute horizontal and vertical gradients around each missing green sample and threshold between them to choose the interpolation direction that runs along, rather than across, a detected edge, directly targeting the zippering problem. Directional filtering with a posteriori decision (DDFAPD, Menon, Andriani and Calvagno, 2007) refined this idea further: two directionally interpolated estimates of the green channel are computed (horizontal and vertical), and a per-pixel decision rule — based only on the reconstructed green channel, rather than requiring full colour images to be computed and compared for both directions — selects the more reliable estimate before red and blue are reconstructed from it. DDFAPD and its variants remained among the strongest model-based baselines through the early 2010s and are still used as reference points in learned-method papers.

2.3 Frequency-Domain, Projection and Regularisation-Based Methods

A separate family of methods reformulates demosaicking as a signal-reconstruction or optimisation problem rather than a purely local interpolation rule. Demosaicking using alternating projections, also referred to as POCS (Projection Onto Convex Sets) demosaicking (Gunturk, Altunbasak and Mersereau, 2002), exploits the high correlation between the high-frequency content of the three colour channels: an initial full-colour estimate is decomposed into sub-bands, the high-frequency detail of the red and blue channels is replaced with that of the green channel, and the result is projected back onto consistency with the originally measured Bayer samples, with the whole procedure iterated for a small number of cycles until convergence. Threshold-based variable number of gradients (VNG) interpolation computes a larger set of directional gradients within a 5×5 neighbourhood and adaptively selects the subset of directions that are most similar to the centre pixel, weighting the corresponding neighbours accordingly. Weighted adaptive colour-plane interpolation (WACPI) and pixel-value data-dependent triangulation (DDT) similarly combine multiple directional estimates — the latter by triangulating each 2×2 pixel block along the diagonal that does not cross a detected outlier pixel, thereby avoiding blending across a likely edge. Block-matching-based interpolation (Acharya and Tsai, 1999) introduces a "colour gravity" similarity metric between candidate neighbourhoods to select the best-matching colour block rather than relying purely on distance-based weighting. Primary-consistent soft-decision demosaicking (Wu and Zhang, 2004) and demosaicking by successive approximation (Li, 2005) push the optimisation view further, respectively combining directionally interpolated estimates via a learned soft classifier and iteratively refining colour-difference estimates through a successive-approximation scheme; the latter remained, at the time of its publication, one of the strongest-performing model-based methods by mean-squared-error on standard test sets.

Table 1. Principal classical, model-based demosaicking algorithms.

Algorithm	Year / Origin	Core Idea	Principal Limitation
Nearest-Neighbour Replication	1987 (Ozawa)	Copies nearest same-colour pixel value	Severe blockiness, no edge awareness

Algorithm	Year / Origin	Core Idea	Principal Limitation
Bilinear Interpolation	1997 (Hamilton & Adams)	Averages nearest same-colour neighbours per channel	Zippering at edges; independent-channel processing
Smooth Hue Transition (Linear / Log)	1987 (Cok)	Interpolates green first, then constrains R/B via hue ratio	Limited gain over bilinear on fine texture
Median-Based Interpolation	1988 (Freeman)	Median-filters colour-difference planes after bilinear pass	Can blur fine chromatic detail
Edge-Sensing Interpolation I / II	1994–1995 (Laroche & Prescott; Hibbard)	Chooses interpolation direction via local gradient thresholding	Threshold sensitivity; local gradients only
Block-Matching Interpolation	1999 (Acharya & Tsai)	Selects neighbouring block via "colour gravity" similarity	No smooth hue constraint; colour bleeding on saturated regions
Threshold-Based Variable Number of Gradients	1999 (Chang et al.)	Adaptively weights a variable subset of directional gradients	Higher computational cost than fixed-neighbourhood methods
Alternating Projections / POCS	2002 (Gunturk et al.)	Iteratively enforces cross-channel high-frequency consistency	Iterative cost; convergence sensitivity
Data-Dependent Triangulation (DDT)	2003 (Su & Willis)	Triangulates each pixel block along the non-outlier diagonal	Sensitive to outlier misclassification
High-Quality Linear Interpolation	2004 (Malvar, He & Cutler)	Bilinear estimate plus gradient-correction gain term	Fixed empirical gain parameters
Primary-Consistent Soft-Decision Demosaicking	2004 (Wu & Zhang)	Combines directional classifiers via learned soft decision	Requires careful weight calibration
Demosaicking by Successive Approximation	2005 (Xin Li)	Iterative colour-difference refinement via decimated signal reconstruction	Weaker performance at well-defined edges
Directional Filtering + A Posteriori Decision (DDFAPD)	2007 (Menon, Andriani & Calvagno)	Green-only directional decision drives red/blue reconstruction	Still outperformed by learned methods on complex texture

3. THE DEEP LEARNING PARADIGM SHIFT

Beginning around 2016, demosaicking research shifted decisively from hand-designed interpolation rules toward learned, data-driven models. The central premise is simple: rather than encoding assumptions about edge orientation and spectral correlation as explicit rules, a sufficiently expressive function — typically a convolutional or attention-based neural network — can be trained end-to-end on large datasets of mosaicked and ground-truth image pairs to learn these regularities directly, and in doing so can capture far more complex, non-linear and context-dependent statistics of natural images than any hand-crafted rule set. This section surveys the resulting literature across nine overlapping architectural and problem-formulation sub-paradigms.

3.1 Early Convolutional Networks

The paper widely credited with establishing deep learning as the dominant paradigm for demosaicking is Gharbi et al.'s "Deep Joint Demosaicking and Denoising" (2016), commonly referred to as DeepJoint. The authors trained a relatively compact convolutional network jointly on demosaicking and denoising, using a purpose-built training set that deliberately over-represented the hard cases — fine textures, saturated regions and diagonal edges — on which classical algorithms perform worst, rather than relying on a generic natural-image corpus. This combination of a targeted, hard-example-mined dataset with an end-to-end trainable network produced a substantial jump in reported PSNR over the best contemporaneous classical baselines, and established the joint demosaicking-denoising formulation, rather than noise-free demosaicking in isolation, as the more practically relevant version of the task, since real camera raw data is never noise-free. Related early CNN approaches, including Deep Camera's fully convolutional ISP network, similarly demonstrated that a single network trained to map raw mosaic to final image could match or outperform hand-engineered multi-stage ISP pipelines across demosaicking, denoising and defective-pixel correction simultaneously.

3.2 Residual, Iterative and Cascade Architectures

A second wave of work reformulated demosaicking as an iterative refinement problem, borrowing ideas from image-restoration research more broadly. Kokkinos and Lefkimmiatis's "Deep Image Demosaicking using a Cascade of Convolutional Residual Denoising Networks" trains a sequence of residual denoising sub-networks, each of which refines the output of the previous stage, and demonstrated that this iterative residual formulation could match state-of-the-art accuracy while using dramatically fewer parameters than earlier single-pass CNNs — on the order of a few hundred thousand parameters compared with over half a million for DeepJoint and well over a million for some contemporaneous alternatives such as Henz et al.'s network. This parameter efficiency matters directly for on-device deployment, where memory bandwidth and model size are often tighter constraints than raw compute. MMNet and related iterative joint demosaicking-denoising networks extended this idea with model-driven, unrolled-optimisation-style architectures that alternate between a data-consistency step (enforcing agreement with the measured mosaic) and a learned denoising/refinement step, echoing the classical POCS formulation but replacing the hand-designed projection with a learned one.

3.3 Attention-Based and Transformer Architectures

Following the broader adoption of self-attention and Vision Transformer architectures across low-level vision from roughly 2021 onward, several demosaicking-specific Transformer variants emerged. RSTCANet applies a residual Swin-Transformer channel-attention design to the demosaicking task and has been reported, across several independent follow-up papers, to be among the strongest performing methods on the standard Kodak and McMaster benchmarks, outperforming both classical baselines and earlier CNN-based methods such as DeepDemosaick by several decibels of PSNR (see Section 6 for consolidated figures). More broadly, general-purpose image-restoration Transformers such as Restormer have been adapted or used as backbones for demosaicking and joint demosaicking-denoising pipelines, exploiting their ability to model long-range spatial dependencies — useful for propagating colour and texture information across larger image regions than the local receptive fields of purely convolutional networks. DemosaicFormer, introduced for the CVPR 2024 MIPI (Mobile Intelligent Photography and Imaging) challenge, adopts a coarse-to-fine Transformer design specifically targeting the HybridEVS Quad-Bayer sensor pattern, reflecting the field's increasing focus on non-standard CFA layouts.

3.4 Implicit Neural Representations

A distinct and more recent line of work reframes demosaicking using implicit neural representations (INRs), in which a multilayer perceptron is trained to map continuous spatial coordinates directly to colour values, conditioned on the observed mosaic, rather than operating on a fixed discrete convolutional grid. NeRD (Neural field-based Demosaicking) demonstrated that this formulation can outperform CNN-based methods such as DeepDemosaick while remaining competitive with, though slightly behind, the strongest Transformer-based methods such as RSTCANet on both the Kodak and McMaster benchmarks. A related 2025 study from the Czech Academy of Sciences explored INR-based demosaicking explicitly for out-of-distribution robustness, noting that conventional pre-trained deep networks such as RSTCANet and DeepDemosaick, while highly accurate on data resembling their training distribution, generalise poorly to raw images containing blur or noise characteristics not represented during training — a limitation that per-image, optimisation-based INR approaches are

structurally less susceptible to, since they fit a fresh representation to each input rather than relying purely on a fixed, pre-trained mapping.

3.5 State-Space (Mamba) and Hybrid Efficient Architectures

The most recent architectural family to be applied to demosaicking, emerging through 2024 and 2025, is based on selective state-space models, popularised in low-level vision under the "Mamba" family of architectures, which offer linear-time sequence modelling as an alternative to the quadratic-cost self-attention used in Transformers. Retinex-RAWMamba (2024) bridges demosaicking and denoising for low-light raw image enhancement using a Mamba-based backbone. The Binarized Mamba-Transformer (BMTNet), presented at CVPR 2025, targets the specific and demanding case of lightweight Quad-Bayer HybridEVS demosaicing: it combines a hybrid Mamba/Swin-Transformer architecture with a binarized semantic Mamba module in which the internal projection layers are binarized for computational efficiency while the core selective-scan computation is retained in full precision, and reports large reductions in memory consumption relative to comparable full-precision Transformer baselines while remaining competitive in reconstruction accuracy — a design point aimed squarely at mobile and embedded deployment. A closely related efficiency-oriented direction is represented by "Efficient Deep Demosaicing with Spatially Downsampled Isotropic Networks" (2026), which restructures the network to operate on a spatially downsampled representation while preserving an isotropic (uniform-resolution) internal architecture, directly targeting inference latency and memory bandwidth on resource-constrained hardware.

3.6 Generative and Deep-Prior Approaches

A smaller but conceptually distinct body of work applies generative modelling ideas to demosaicking. Deep Image Prior (DIP)-based demosaicking exploits the observation, originally made in the broader image-restoration literature, that the structure of a convolutional network architecture alone — without any pre-training on external data — acts as an implicit prior favouring natural-image statistics; a randomly initialised network is optimised, per image, to reproduce the observed mosaic samples, and the interpolated result at convergence is taken as the demosaicked output. Reported results place DIP-based demosaicking close to, but for most test images not exceeding, the strongest directional classical methods such as Menon et al.'s DDFAPD, reflecting the inherent difficulty of matching fully supervised methods without access to any external training data. Plug-and-play (PnP) approaches occupy a related middle ground: rather than training a demosaicking-specific network end-to-end, a generic pre-trained denoiser — such as DRUNet — is inserted as a modular prior into an iterative optimisation scheme (for example, ADMM), alternating between enforcing consistency with the measured mosaic and applying the learned denoiser as a regulariser. Preconditioned PnP-ADMM variants using a demosaicking-specialised DRUNet prior (DRUNet-dem) have been reported to match or exceed dedicated end-to-end networks such as Gharbi et al.'s DeepJoint on both the Kodak and McMaster benchmarks (see Section 6), while retaining the modularity and interpretability advantages of the optimisation-based framework.

3.7 Joint Demosaicing-Denoising and Non-Bayer CFA Models

As noted in Section 1.2, the proliferation of Quad-Bayer, Nona-Bayer and HybridEVS sensor layouts in smartphone cameras has made pattern-specific demosaicking an increasingly central research problem in its own right, and joint demosaicking-denoising — rather than noise-free demosaicking — the practically dominant formulation, since real sensor readouts are never noise-free and sequential denoise-then-demosaic or demosaic-then-denoise pipelines are known to compound each other's errors. DRNet, presented at AAAI 2024, introduces a dual-encoder network with joint residual learning specifically for Quad-Bayer joint demosaicing and denoising, using a self-adaptive encoder and a joint residual block that independently estimates the demosaicing and denoising residual components. "Beyond Joint Demosaicking and Denoising", an ICPR 2022 paper, proposes a full image-processing pipeline tailored to pixel-binning sensors and benchmarks several CNN-based joint demosaicing-denoising methods (DeepJoint, Kokkinos et al., DeepISP and others) specifically on Quad-Bayer data. Most notably, "Examining Joint Demosaicing and Denoising for Single-, Quad-, and Nona-Bayer Patterns" (Tedla, Punnappurath, Zhao and Brown, presented at ICCP 2025) directly addresses the operational inefficiency of training a separate specialised network per sensor pattern — a real deployment burden for smartphone manufacturers now shipping multiple camera modules with different CFA layouts on a single device. The authors show that concatenating a three-channel mosaic-

layout embedding to the network input and training a single unified architecture across all three pattern types outperforms existing pattern-specific Quad-Bayer and Nona-Bayer models and achieves accuracy comparable to dedicated Single-Bayer models, while additionally introducing a masking strategy that improves robustness to dead or defective sensor pixels, a practical failure mode largely overlooked by prior AI-based demosaicking research. As part of this effort the authors also released a new, purpose-captured RAW demosaicing dataset of 638 images spanning challenging real-world scene content.

3.8 Lightweight and Mobile-Deployment-Oriented Models

Because demosaicking must, in a production camera pipeline, run at or near frame rate on power- and memory-constrained mobile system-on-chip hardware, model efficiency is not a secondary concern but often the primary design constraint. Beyond the binarized and spatially-downsampled architectures discussed in Section 3.5, several lines of work explicitly benchmark accuracy against parameter count and inference latency rather than accuracy alone. Kokkinos and Lefkimmiatis's cascaded residual network, discussed in Section 3.2, remains a frequently cited efficiency reference point precisely because it demonstrated competitive accuracy at a fraction of the parameter budget of contemporaneous alternatives. More recent lightweight designs, including compact Inception-style blocks integrated into two-stage joint demosaicking-denoising training strategies, similarly report that carefully designed compact architectures can rank competitively against much larger networks on standard benchmarks, narrowing, though not eliminating, the accuracy-efficiency trade-off that dominates the mobile deployment literature.

3.9 Summary of the Deep Learning Landscape

Table 2 (Section 6) consolidates representative architectures from each of the sub-paradigms above alongside their reported benchmark performance, illustrating the overall trajectory: from an approximately 4–6 dB PSNR improvement of early CNNs over the strongest classical baselines, through a further 2–4 dB gain from attention and Transformer-based architectures, to the current frontier defined less by incremental PSNR gains on clean, noise-free, single-Bayer benchmarks — which are approaching a point of diminishing returns — and more by generalisation across sensor patterns and noise conditions, and by the efficiency required for real-time, on-device operation.

4. BENCHMARK DATASETS

Comparability across the demosaicking literature is complicated by the fact that different papers train and evaluate on different datasets, different re-mosaicking protocols (synthetic re-mosaicking of already-demosaicked sRGB images versus true linear-raw capture), and different noise conditions, a caveat that is discussed further in Section 6. The datasets below recur most frequently across the literature surveyed in this paper.

The Kodak PhotoCD dataset, comprising 24 uncompressed 768×512 images, is by a wide margin the most commonly used demosaicking benchmark, valued for its long history of consistent use despite being small, dated, and not fully representative of the statistics of modern high-resolution camera images. The McMaster (McM) dataset, comprising 18 cropped high-resolution images with deliberately diverse and challenging colour and texture content, is typically reported alongside Kodak as a complementary, higher-difficulty benchmark. BSD68 and BSD100, drawn from the Berkeley Segmentation Dataset, and Urban100, composed of urban scenes with abundant fine repetitive structure such as building facades particularly prone to aliasing, are widely used in the general image-restoration literature and have been adopted by several demosaicking papers, most notably to stress-test performance on high-frequency, alias-prone content. Gharbi et al. additionally introduced a dedicated, purpose-built "MIT" demosaicking training and evaluation dataset that deliberately over-samples the hard cases — saturated regions, fine texture and diagonal edges — on which earlier methods were known to fail, and several subsequent papers report results on this set alongside Kodak and McMaster. The Waterloo Exploration Database (WED) is used in several joint demosaicking-denoising papers as an additional large-scale natural-image source. The MSR Demosaicking Dataset provides real linear-RGB raw captures (as opposed to synthetically re-mosaicked sRGB images) and is used specifically to validate methods under realistic sensor noise and colour statistics. Finally, reflecting the shift toward non-Bayer sensor layouts, the 2025 SamsungLabs unified-demosaicing dataset (638 RAW images with annotated patches for training, validation and testing) was purpose-

built to support joint demosaicing-denoising research across Single-, Quad- and Nona-Bayer layouts simultaneously.

Table 3. Benchmark datasets used across the demosaicking literature.

Dataset	Size / Content	Typical Use
Kodak PhotoCD	24 images, 768×512	Primary standard benchmark across nearly all eras
McMaster (McM)	18 high-resolution crops, diverse texture/colour	Higher-difficulty complementary benchmark
BSD68 / BSD100	68 / 100 natural images (Berkeley Segmentation Dataset)	General restoration benchmark, adopted for demosaicking
Urban100	100 urban-scene images with fine repetitive structure	Aliasing / high-frequency structure stress test
Gharbi "MIT" Demosaicking Set	Large curated set over-sampling hard cases	Training and evaluation for CNN-based methods
Waterloo Exploration Database (WED)	Large-scale natural image corpus	Large-scale training for joint demosaic-denoise models
MSR Demosaicking Dataset	Real linear-RGB raw captures	Evaluation under realistic sensor noise/colour statistics
SamsungLabs Unified-Demosaicing Set (2025)	638 RAW images, annotated patches	Single-/Quad-/Nona-Bayer joint demosaic-denoise research

5. EVALUATION METRICS

Peak signal-to-noise ratio (PSNR), computed on the reconstructed RGB image against ground truth, is the dominant quantitative metric throughout this literature, owing to its simplicity and long precedent, despite being an imperfect proxy for perceived visual quality. The structural similarity index (SSIM) is almost universally reported alongside PSNR as a complementary metric more sensitive to local structural and contrast distortions than to uniform intensity error. For colour-critical applications, the CIE ΔE2000 colour-difference metric is used in several papers, particularly those concerned with false-colour and colour-fringing artefacts that PSNR, computed globally across all three channels, can under-penalise. Mean squared error (MSE), the un-normalised precursor to PSNR, is still reported in some older and some more recent papers directly, particularly where comparing precise reconstruction error rather than a logarithmic quality score is more informative. Beyond reconstruction accuracy, papers targeting real-time or mobile deployment additionally report parameter count, floating-point operations (FLOPs) or multiply-accumulate operations, peak memory consumption, and measured inference latency on representative hardware, reflecting the field's growing emphasis on the accuracy-efficiency trade-off discussed in Section 3.8.

6. CONSOLIDATED COMPARATIVE ANALYSIS

This section brings together quantitative results reported in the primary sources surveyed in Section 3, in order to give a concrete sense of the magnitude of improvement delivered by successive methodological generations. Two important caveats apply to any such cross-paper comparison and should be borne in mind throughout this section. First, different papers use different training data, different re-mosaicking or true-raw evaluation protocols, and in some cases slightly different crop or border-handling conventions when computing PSNR, so figures compiled from different source papers are indicative of relative trends rather than perfectly controlled comparisons. Second, most of the figures reported here correspond to the noise-free demosaicking setting, which is a useful common ground for comparison across three decades of literature but is, as discussed in Sections 3.7 and 7, an

increasingly unrepresentative simplification of the true, noisy, non-Bayer operating conditions that current research targets.

Table 2 reports average PSNR and SSIM figures on the Kodak and McMaster datasets for a representative cross-section of methods spanning the classical and deep-learning eras, compiled from the NeRD paper (Kerepecky et al. and related implicit-representation literature) and from the preconditioned plug-and-play ADMM study, both of which independently tabulate multiple prior methods on a consistent evaluation protocol within their respective papers.

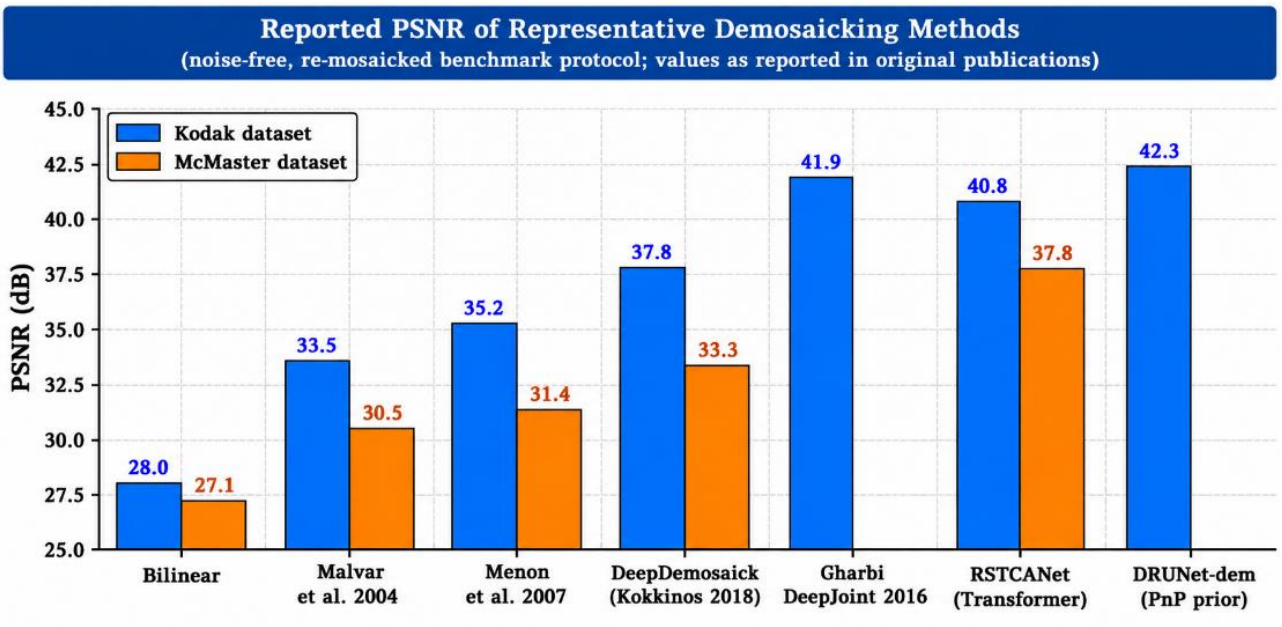

Figure 3. Reported PSNR of representative demosaicking methods on the Kodak and McMaster datasets, compiled from the source papers listed in Table 2.

Table 2. Reported PSNR on Kodak and McMaster datasets across representative methods (noise-free demosaicking; figures as tabulated within the cited source papers).

Method	Category	Kodak PSNR (dB)	McMaster PSNR (dB)	Source
Bilinear Interpolation	Classical	28.01	27.15	NeRD study (2023)
Malvar et al. (2004)	Classical	33.52 / 34.68	30.54 / 34.46	NeRD study; PnP-ADMM study
Menon et al. – DDFAPD (2007)	Classical	35.20	31.40	NeRD study (2023)
DeepDemosaick – Kokkinos & Lefkimmiatis (2018)	CNN (residual cascade)	37.76	33.31	NeRD study (2023)
Gharbi et al. – DeepJoint (2016)	CNN (end-to-end)	41.85	39.13	PnP-ADMM study (2021)
Henz et al. / Kokkinos – MMNet	CNN (iterative / model-driven)	40.26–42.05	36.84–39.50	PnP-ADMM study (2021)

Method	Category	Kodak PSNR (dB)	McMaster PSNR (dB)	Source
RSTCANet	Transformer (Swin channel-attention)	40.84	37.77	NeRD study (2023)
NeRD (Implicit Neural Representation)	Implicit neural representation	39.07	36.18	NeRD study (2023)
Preconditioned PnP-ADMM (DRUNet-dem prior)	Plug-and-play deep prior	42.32–42.37	39.27–39.31	PnP-ADMM study (2021)

7. CROSS-CUTTING CHALLENGES

Several challenges recur across the literature surveyed in this paper and, taken together, define the current research frontier.

7.1 Generalisation Beyond the Training Distribution

Deep demosaicking networks are trained on a fixed corpus of mosaicked-clean image pairs and, as the 2025 implicit-neural-representation study explicitly documents, tend to generalise poorly to raw images exhibiting noise levels, blur characteristics, or sensor-specific colour statistics not well represented during training. This is a meaningful practical limitation: a network trained and validated on Kodak- or McMaster-style clean, re-mosaicked sRGB data may substantially under-perform on genuine sensor raw data captured under different illumination, ISO or optical conditions than its training set.

7.2 Fragmentation Across Sensor Patterns

As discussed in Section 3.7, the historical norm of training a separate, pattern-specific network for each of Single-Bayer, Quad-Bayer, Nona-Bayer and now HybridEVS sensor layouts creates a real deployment burden for camera and smartphone manufacturers who must maintain, validate and run multiple specialised models. Unified multi-pattern architectures, such as the 2025 mosaic-embedding approach, are a promising but still early response to this fragmentation, and broader adoption of similarly unified formulations is likely to be an active research direction through the remainder of the decade.

7.3 The Accuracy-Efficiency Trade-Off

The strongest-performing architectures on standard benchmarks — large Transformer and hybrid state-space networks — are frequently far too computationally expensive for real-time, on-device execution on mobile image signal processors, which must typically demosaic a multi-megapixel frame within tens of milliseconds under strict power budgets. Binarization, spatial downsampling, and compact cascade architectures (Sections 3.5 and 3.8) represent partial responses to this constraint, but a persistent, non-trivial accuracy gap remains between the most efficient deployable models and the most accurate research-grade models, and closing this gap without sacrificing real-time performance remains an open problem.

7.4 Joint Optimisation Versus Modular ISP Design

A recurring architectural question is whether demosaicking should be optimised jointly with denoising, and potentially with other ISP stages such as white balance and tone mapping, within a single end-to-end trainable pipeline, or kept as a modular, independently optimised stage within a conventional multi-stage ISP. Joint formulations (Sections 3.1 and 3.7) generally report accuracy gains over strictly sequential pipelines, since sequential demosaic-then-denoise or denoise-then-demosaic processing is known to compound each stage's characteristic errors; however, fully joint, end-to-end ISP networks sacrifice the interpretability, independent testability and incremental upgradability that a modular pipeline affords, which remains a practically important consideration for camera manufacturers maintaining large, heterogeneous product lines.

7.5 Evaluation Protocol Inconsistency

As noted in Section 6, the field lacks a fully standardised evaluation protocol: some papers evaluate on synthetically re-mosaicked sRGB images (which do not fully capture true sensor raw statistics), others on genuine linear-RGB raw captures; noise levels, when included, vary across studies; and border-cropping conventions for PSNR computation are not uniformly reported. This inconsistency, while not unique to demosaicking within the broader image-restoration literature, makes precise cross-paper ranking difficult and is a contributing reason why this survey emphasises reported trends over any claim of exact numerical ranking beyond what individual source papers directly compare.

8. FUTURE RESEARCH DIRECTIONS

Based on the trajectory of the literature reviewed above, several directions appear particularly likely to shape demosaicking research over the coming several years.

Unified, pattern-agnostic architectures. Extending the mosaic-embedding approach demonstrated for Single-, Quad- and Nona-Bayer patterns toward a genuinely pattern-agnostic architecture capable of adapting to arbitrary or even previously unseen CFA layouts — including further HybridEVS and application-specific sensor designs — without per-pattern retraining, would directly address the deployment fragmentation discussed in Section 7.2.

Self-supervised and unsupervised training regimes. Because paired mosaicked/ground-truth training data is comparatively easy to synthesise from existing clean images but true sensor-raw and ground-truth pairs are harder to obtain at scale, self-supervised formulations that can learn directly from unpaired or weakly paired raw camera data — potentially exploiting techniques such as noise2noise-style training or cycle-consistency between re-mosaicked and demosaicked representations — are a promising route to improving real-world generalisation without requiring large curated paired datasets.

Diffusion and other generative priors. Diffusion models have recently demonstrated strong performance as learned priors across a range of inverse imaging problems, including super-resolution and inpainting, but remain comparatively underexplored specifically for demosaicking; given the structural similarity between demosaicking and other under-determined linear inverse problems, diffusion-based or score-based generative priors, potentially integrated within a plug-and-play framework similar to the DRUNet-based approaches discussed in Section 3.6, represent a natural and likely fruitful extension.

Efficient state-space and hybrid architectures for on-device inference. Building on the binarized Mamba-Transformer and spatially downsampled isotropic network directions introduced in 2025 and 2026, further work on quantisation-aware training, structured pruning, and hardware-co-designed state-space architectures is likely to continue narrowing the accuracy-efficiency gap identified in Section 7.3.

Joint, fully end-to-end neural ISP pipelines. Extending joint demosaicing-denoising formulations to additionally incorporate white balance, tone mapping and even downstream task-specific objectives (for example, optimising the full pipeline directly for object-detection accuracy rather than pixel-level fidelity) within a single differentiable pipeline is an active direction, balanced against the modularity trade-offs discussed in Section 7.4.

Standardised, sensor-realistic benchmarks. Given the evaluation inconsistencies discussed in Section 7.5, the field would likely benefit from a broader, community-standard benchmark suite that combines multiple genuine sensor patterns, realistic noise models calibrated against real camera sensors, and a fixed, published evaluation protocol, analogous to the role played by standardised benchmarks in other areas of computer vision.

9. CONCLUSION

Colour filter array demosaicking has progressed, over nearly four decades, from simple nearest-neighbour and bilinear interpolation through increasingly sophisticated edge-aware, frequency-domain and optimisation-based classical algorithms, and finally into a deep-learning era that has itself already passed through several distinct architectural generations: early convolutional networks, residual and iterative cascades, attention and Transformer-based models, implicit neural representations, and, most recently, state-space and hybrid binarized architectures. Reported benchmark results compiled in this survey show reconstruction quality improving by roughly eight to fourteen decibels of PSNR on

standard test sets across this progression, a substantial and practically significant gain. At the same time, the field's centre of gravity has visibly shifted away from incremental accuracy gains on clean, noise-free, single-Bayer benchmarks — which appear to be approaching a point of diminishing returns — and toward the harder, more practically consequential problems of joint demosaicing-denoising under realistic sensor noise, generalisation across the rapidly diversifying landscape of Quad-Bayer, Nona-Bayer and HybridEVS sensor layouts, and efficient, real-time inference on power-constrained mobile hardware. Addressing these problems in a unified, generalisable and deployable manner, rather than through ever more specialised, pattern- and dataset-specific models, is likely to define the next phase of demosaicking research.

DECLARATIONS:

Acknowledgments	:	Not applicable.
Conflict of Interest	:	The author declares that there is no actual or potential conflict of interest about this article.
Consent to Publish	:	The authors agree to publish the paper in the Global Research Journal of Social Sciences and Management.
Ethical Approval	:	Not applicable.
Funding	:	Author claims no funding was received.
Author Contribution	:	Both the authors confirms their responsibility for the study, conception, design, data collection, and manuscript preparation.
Data Availability Statement	:	The data presented in this study are available upon request from the corresponding author.
Use of Generative Artificial Intelligence	:	AI-assisted tools have been used to improve the images, language editing or grammatical clarity, the authors take full responsibility for the content of the final version.

REFERENCES

[1] Bayer, B. E., "Colour Imaging Array," U.S. Patent 3,971,065, Eastman Kodak Company, 1976.

[2] Ozawa, N., "Chrominance Signal Interpolation Device for a Colour Camera," U.S. Patent 4,716,455, 1987.

[3] Cok, D. R., "Signal Processing Method and Apparatus for Producing Interpolated Chrominance Values in a Sampled Colour Image Signal," U.S. Patent 4,642,678, Eastman Kodak Company, 1987.

[4] Freeman, W. T., "Median Filter for Reconstructing Missing Colour Samples," U.S. Patent 4,724,395, 1988.

[5] Laroche, C. A., Prescott, M. A., "Apparatus and Method for Adaptively Interpolating a Full Colour Image Utilizing Chrominance Gradients," U.S. Patent 5,373,322, Eastman Kodak Company, 1994.

[6] Hibbard, R. H., "Apparatus and Method for Adaptively Interpolating a Full Colour Image Using Luminance Gradients," U.S. Patent 5,382,976, Eastman Kodak Company, 1995.

[7] Hamilton, J., Adams, J., "Adaptive Colour Plane Interpolation in Single Sensor Colour Electronic Camera," U.S. Patent 5,629,734, 1997.

[8] Acharya, T., Tsai, P. S., "A New Block Matching Based Colour Interpolation Algorithm," Proc. SPIE Electronic Imaging: Science and Technology, vol. 3648, pp. 60–65, 1999.

[9] Chang, E. et al., "Colour Filter Array Recovery Using a Threshold-Based Variable Number of Gradients," Proc. SPIE, 1999.

[10] Gunturk, B. K., Altunbasak, Y., Mersereau, R. M., "Colour Plane Interpolation Using Alternating Projections," IEEE Transactions on Image Processing, vol. 11, no. 9, pp. 997–1013, 2002.

[11] Su, D., Willis, P., "Demosaicing of Colour Images Using Pixel Level Data-Dependent Triangulation," Theory and Practice of Computer Graphics, IEEE Proceedings, pp. 16–23, 2003.

[12] Malvar, H. S., He, L.-W., Cutler, R., "High-Quality Linear Interpolation for Demosaicing of Bayer-Patterned Colour Images," Proc. IEEE ICASSP, 2004.

[13] Wu, X., Zhang, N., "Primary-Consistent Soft-Decision Colour Demosaicking for Digital Cameras," IEEE Transactions on Image Processing, vol. 13, no. 9, pp. 1263–1274, 2004.
[14] Xin Li, Orchard, M. T., "New Edge-Directed Interpolation," IEEE Transactions on Image Processing, vol. 10, no. 10, 2001.
[15] Xin Li, "Demosaicing by Successive Approximation," IEEE Transactions on Image Processing, vol. 14, no. 3, pp. 370–379, 2005.
[16] Menon, D., Andriani, S., Calvagno, G., "Demosaicing With Directional Filtering and a Posteriori Decision," IEEE Transactions on Image Processing, vol. 16, no. 1, pp. 132–141, 2007.
[17] Gharbi, M., Chaurasia, G., Paris, S., Durand, F., "Deep Joint Demosaicking and Denoising," ACM Transactions on Graphics (Proc. SIGGRAPH Asia), vol. 35, no. 6, 2016.
[18] Kokkinos, F., Lefkimmiatis, S., "Deep Image Demosaicking Using a Cascade of Convolutional Residual Denoising Networks," Proc. European Conference on Computer Vision (ECCV), 2018.
[19] Xing, W. et al., "RSTCANet: Residual Swin Transformer Channel Attention Network for Image Demosaicing," IEEE Signal Processing Letters, 2022 (as cited in subsequent demosaicking literature).
[20] Zamir, S. W., Arora, A., Khan, S., Hayat, M., Khan, F. S., Yang, M.-H., "Restormer: Efficient Transformer for High-Resolution Image Restoration," Proc. IEEE/CVF CVPR, 2022.
[21] Wu, Y., Fan, Z., Chu, X. et al., "DemosaicFormer: Coarse-to-Fine Demosaicing Network for HybridEVS Camera," Proc. IEEE/CVF CVPR Workshops, 2024.
[22] Kerepecký, T., Šroubek, F., Flusser, J., "NeRD: Neural Field-Based Demosaicking," arXiv:2304.06566, 2023.
[23] Kerepecký, T., Šroubek, F., Flusser, J., "Implicit Neural Representations for Robust Joint Demosaicking and Denoising," Digital Signal Processing, vol. 159, article 105022, 2025.
[24] Balušík, P., "Image Demosaicing Using Deep Image Prior," Signal Processing Laboratory, Brno University of Technology (supervised by P. Rajmic), 2022.
[25] Zhang, K., Li, Y., Zuo, W., Zhang, L., Van Gool, L., Timofte, R., "Plug-and-Play Image Restoration with Deep Denoiser Prior," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 44, no. 10, pp. 6360–6376, 2021.
[26] Anonymous et al., "Preconditioned Plug-and-Play ADMM with Locally Adjustable Denoiser for Image Restoration," arXiv:2110.00493, 2021.
[27] Sharif, S. M. A., Naqvi, R. A., Biswas, M., "Beyond Joint Demosaicking and Denoising: An Image Processing Pipeline for a Pixel-Bin Image Sensor," Proc. International Conference on Pattern Recognition (ICPR), 2022.
[28] Anonymous et al., "Quad Bayer Joint Demosaicing and Denoising Based on Dual Encoder Network with Joint Residual Learning (DRNet)," Proc. AAAI Conference on Artificial Intelligence, 2024.
[29] Anonymous et al., "Joint Demosaicking and Denoising Benefits from a Two-Stage Training Strategy," arXiv:2009.06205, 2020.
[30] Kim, J., Kim, M. H., "Joint Demosaicing and Deghosting of Time-Varying Exposures for Single-Shot HDR Imaging," Proc. IEEE/CVF ICCV, 2023.
[31] Tedla, S., Punnappurath, A., Zhao, L., Brown, M. S., "Examining Joint Demosaicing and Denoising for Single-, Quad-, and Nona-Bayer Patterns," Proc. IEEE International Conference on Computational Photography (ICCP), 2025.
[32] Zhou et al., "Binarized Mamba-Transformer for Lightweight Quad Bayer HybridEVS Demosaicing," Proc. IEEE/CVF CVPR, 2025.
[33] Anonymous et al., "Retinex-RAWMamba: Bridging Demosaicing and Denoising for Low-Light RAW Image Enhancement," arXiv, 2024.
[34] Anonymous et al., "Efficient Deep Demosaicing with Spatially Downsampled Isotropic Networks," arXiv:2601.00703, 2026.
[35] Anonymous et al., "Event Camera Demosaicing via Swin Transformer and Pixel-Focus Loss," arXiv:2404.02731, 2024 (CVPR MIPI Challenge).
[36] Anonymous et al., "Adaptive Cross-Layer Attention for Image Restoration," arXiv:2203.03619, 2022.

[37] Tan, D. et al., "Deep Camera: A Fully Convolutional Neural Network for Image Signal Processing," arXiv:1908.09191, 2019.
[38] Guo, T., Zhou, Q., Gao, M., Jeon, G., Camacho, D., "Demosaicking Algorithm Using Deep Residual Convolutional Network," International Journal of Interactive Multimedia and Artificial Intelligence, vol. 9, no. 7, pp. 6–15, 2026.

Authors' Profiles

Mr. C. Raja Rao obtained his B.E in ECE from Andhra University in 1995 and M.Tech in ECE from JNTU Hyderabad in 2001. At present he is pursuing his Ph.D Degree in Electrical Engineering from Maulana Abul Kalam Azad University of Technology, formerly known as West Bengal University of Technology, Haringhata, West Bengal, India. He started his career in 1999 as Assistant Professor of Electronics and Communication Engineering Department with SVITS Mahbubnagar. He worked in various Engineering colleges and continued in teaching up to Dec 2009. In December 2009 he has been appointed as Deputy Director in Board of Practical Training (Eastern Region) Kolkata (An Autonomous body under Ministry of Education Government of India). He has over 10 years of teaching experience and about 16 years of experience in implementation of Apprentices Act 1961. He is also a life member of ISTE, IETE and IE(I). His research interest includes Signal and Image Processing, Low power VLSI design, Vocational and Technical Education and Training.

Dr. Soumitra Kumar Mandal has obtained his B.E. from Bengal Engineering College (Now IIEST), Shibpur, M.Tech from Institute of Technology, Banaras Hindu University, Varanasi and Ph.D. from Punjab University, Chandigarh all in Electrical Engineering. He started his career as Lecturer at SSGM Engineering College, Shegaon, and then moved to Punjab Engineering College, Chandigarh. In February, 2004, he has been appointed as Assistant Professor of Electrical Engineering in National Institute of Technical Teachers' Training and Research (NITTTR), Kolkata. He is now serving as Professor of Electrical Engineering in the same institute. Throughout his academic career, he has published about 45 research papers in National and International Journals and presented many papers in conferences. He has also published 8 Text books for undergraduate and Post Graduate Students of Electrical Engineering. His research interests include Microprocessor and Microcontroller based System Design, Embedded System Design, Computer Controlled Drives, Neuro-fuzzy Computing, Signal Processing and VLSI design. He is also a life member of ISTE and a member of IE.