

An End-to-End Real-World Camera Imaging Pipeline

Kepeng Xu
Xidian University
Xi'an, China
kepengu11@gmail.com

Zijia Ma
Xidian University
Xi'an, China

Li Xu
Xidian University
Xi'an, China
icecherylxuli@gmail.com

Gang He*
Xidian University
Xi'an, China
ghe@xidian.edu.cn

Yunsong Li
Xidian University
Xi'an, China
ysli@mail.xidian.edu.cn

Wenxin Yu
Southwest University of Science and
Technology
Mianyang, China

Taichu Han
Novastar Tech Co., Ltd.
Xi'an, China

Cheng Yang
Novastar Tech Co., Ltd.
Xi'an, China

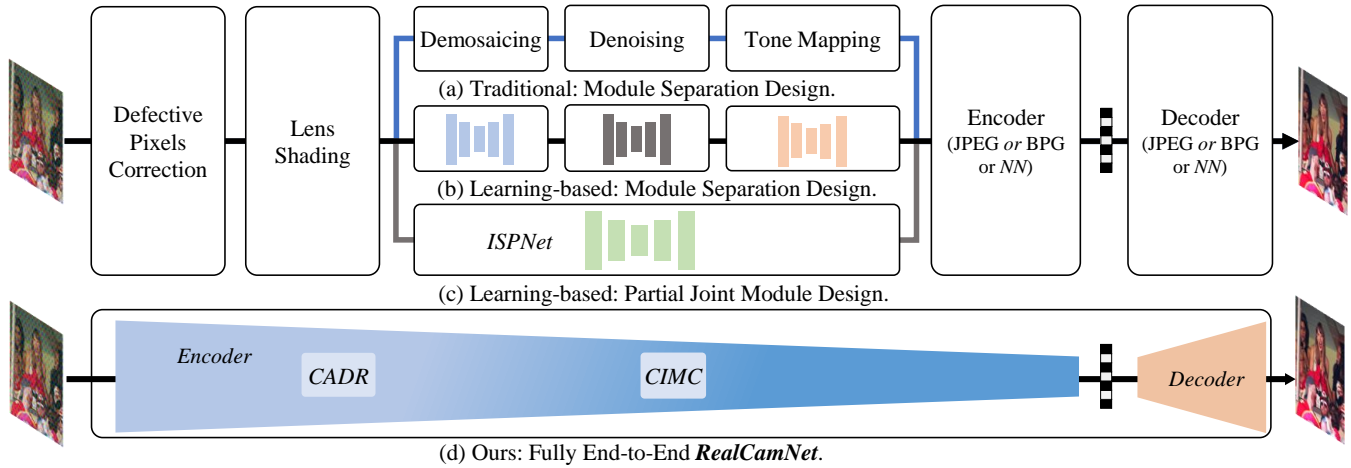


Figure 1: (a) Traditional Design: Implements each step of Image Signal Processing (ISP) separately using conventional algorithms. (b) Learning-based Separated Design: Employs neural networks to individually address each phase of the ISP process. (c) Learning-based Partial Joint Design: Develop an ISPNet to unify operations such as demosaicing, and tone mapping. (d) RealCamNet: We propose an end-to-end camera imaging framework that categorizes ISP operations into coordinate-independent and coordinate-dependent groups. CIMC and CADR are designed to perform tasks like demosaicing and image feature compression and to restore coordinate-dependent image distortions (e.g., vignetting, dark shadows), respectively.

ABSTRACT

Recent advances in neural camera imaging pipelines have demonstrated notable progress. Nevertheless, the real-world imaging

pipeline still faces challenges including the lack of joint optimization in system components, computational redundancies, and optical distortions such as lens shading. In light of this, we propose an end-to-end camera imaging pipeline (RealCamNet) to enhance real-world camera imaging performance. Our methodology diverges from conventional, fragmented multi-stage image signal processing towards end-to-end architecture. This architecture facilitates joint optimization across the full pipeline and the restoration of coordinate-biased distortions. RealCamNet is designed for high-quality conversion from RAW to RGB and compact image compression. Specifically, we deeply analyze coordinate-dependent optical distortions, e.g., vignetting and dark shading, and design a novel Coordinate-Aware Distortion Restoration (CADR) module to restore coordinate-biased distortions. Furthermore, we propose a Coordinate-Independent Mapping Compression (CIMC) module to

*Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM'24, October 28 - November 1, 2024, Melbourne, Australia.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/18/06

<https://doi.org/XXXXXXXX.XXXXXXX>

implement tone mapping and redundant information compression. Existing datasets suffer from misalignment and overly idealized conditions, making them inadequate for training real-world imaging pipelines. Therefore, we collected a real-world imaging dataset. Experiment results show that RealCamNet achieves the best rate-distortion performance with lower inference latency.

KEYWORDS

Camera imaging, deep neural network, image compression, image signal processing

ACM Reference Format:

Kepeng Xu, Zijia Ma, Li Xu, Gang He, Yunsong Li, Wenxin Yu, Taichu Han, and Cheng Yang. 2018. An End-to-End Real-World Camera Imaging Pipeline. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (MM'24)Proceedings of the 32nd ACM International Conference on Multimedia (MM'24), October 28-November 1, 2024, Melbourne, Australia*. ACM, New York, NY, USA, 10 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

Efficient imaging and compression technologies are paramount to the internet and multimedia industries. With the exponential increase in digital content, techniques that reduce file sizes while maintaining image quality have become crucial. These advancements not only enhance the user experience by ensuring fast image and video loading but also help save storage space and reduce data transmission costs. Moreover, such technologies are indispensable for supporting advanced applications like autonomous driving and remote sensing, driving technological innovation, and meeting the growing demands for multimedia processing. Therefore, developing more efficient imaging and compression methods is vital for propelling industry progress and meeting the needs of modern technology.

In traditional real-world camera imaging pipelines, complex and proprietary hardware processes are employed for image signal processing (ISP), encompassing steps such as denoising, demosaicing, tone mapping, and image compression, as shown in Fig.1 (a). The pipeline's step-by-step design, where each process is designed separately, leads to error accumulation and prevents the system from achieving an optimal state.

Most traditional methods derive solutions at each step of the ISP pipeline using heuristic approaches[14, 29, 40], thus requiring the adjustment of numerous parameters. In addition, errors will be accumulated in the ISP processing flow, affecting imaging quality.

Deep neural networks have developed rapidly in the past decade, and are used in image classification[34, 36], object detection[27, 45], image and video enhancement[2, 26, 39, 44, 46], natural language processing[9, 24, 37] and other fields have played an important role. Researchers have also proposed a series of image signal processing methods based on neural networks. In recent years, researchers[15, 23] have attempted to construct camera imaging systems based on deep neural networks, designing models for denoising, demosaicing, tone mapping, and compression, significantly enhancing overall system performance. However, constructing separate neural networks for each function (e.g., denoising, demosaicing) prevents the system from being jointly optimized and introduces computational redundancy, as illustrated in Fig.1 (b).



Figure 2: Coordinate-dependent distortion. The left side is the captured image, and the right side is the vignetting distribution map. The middle of the picture on the left is brighter.

Therefore, some methods[18, 19, 48] have emerged to try to build a single neural network to complete the functions required by ISP. This is a good idea, but because the situation considered is too simple, the designed ISPNet cannot be realistically applied.

To address these issues, we propose the RealCamNet framework, which integrates the real camera image signal processing and image compression processes into an end-to-end deep neural network framework. This unified approach not only reduces computational redundancy but also enhances both image quality and compression efficiency through the joint optimization of network parameters, a feat unattainable by previously isolated design methods.

Our starting point was to construct a neural network-based real-world camera pipeline that simulates common functions found in real-world applications, building a reliable and practical end-to-end imaging system. To this end, we analyzed the principles and flaws in each step of the imaging and compression process.

Our RealCamNet's overall framework is shown in Fig.1 (d), where we designed an end-to-end deep neural network to implement the RAW->Bitstream->RGB imaging compression process. The past cascaded framework depicted in Fig.1 (a) and Fig.1 (b) leads to problems such as error accumulation from each module and potentially suboptimal global results. Fig.1 (c) constructing a neural network directly implements operations such as demosaicing and tone mapping, but the overall joint optimization is still not achieved.

In our RealCamNet, we designed global and local perception imaging pipeline modules to simulate processes in the imaging pipeline, such as global and local tone mapping, denoising, demosaicing, and image feature compression.

Our investigation into camera imaging technologies has identified digital signal distortions stemming from optical system imperfections. In optical camera imaging, Coordinate-dependent distortions significantly impact image quality due to the inherent characteristics of the optical system and sensor. These include Coordinate-dependent noise resulting from sensor sensitivity variations or uneven optical system illumination, and vignetting, a reduction in edge brightness caused by lens design and light entry angles. As well as dark shading caused by problems such as uneven heating of CMOS due to the superposition of components. The vignetting and dark shading is shown in Fig.2. To mitigate these distortions and improve image quality, sophisticated image processing is required. We introduce a streamlined, effective coordinate-aware distortion correction module designed to identify and rectify various Coordinate-biased distortions, enhancing visual perception and imaging quality.

Previous RAW-RGB datasets either have too simple ISP processes (only include a few processes such as demosaicing), or have misalignment issues and cannot be used to train camera imaging pipelines that can be used in real scenes. To train our RealCamNet, we construct a new RAW-RGB dataset to directly learn the complex ISP process of real cameras.

In summary, our contributions are fourfold:

- We constructed a deep neural network imaging system and demonstrated the performance improvements from end-to-end joint optimization.
- We deeply analyze coordinate-dependent optical distortions, e.g., vignetting and dark shading, and design a novel Coordinate-Aware Distortion Restoration (CADR) module to repair coordinate-biased distortions.
- We designed a Coordinate-Independent Mapping Compression (CIMC) module to implement global and local tone mapping, denoising, demosaicing, and image feature compression functions, thereby reducing computational redundancy.
- We built a new real-world imaging dataset, providing a benchmark for the unified end-to-end neural imaging pipeline.

2 PRELIMINARY

2.1 Traditional Image Signal Processing

Traditionally, the Image Signal Processing (ISP) is responsible for reconstructing RGB images from RAW captures. In conventional camera pipelines, complex and proprietary hardware processes are employed for image signal processing. This process encompasses several steps, including denoising, demosaicing, defect pixel correction, and tone mapping[6, 11, 13, 28, 29, 41], among others. Each module requires individual tuning and optimization, with consideration for the adjustment of cascading errors.

2.2 Learning-based Image Signal Processing

Advancements in deep learning have led researchers to enhance ISP pipelines via neural networks. Localized network solutions segment functions such as tone mapping[35] and denoising[16, 23], facilitating module decoupling yet complicating the ISP's holistic optimization. Comprehensive neural network approaches replace the entire ISP pipeline, streamlining computations and enabling system-wide optimization[10, 20, 21, 30, 49].

2.3 Learning-based Image Compression

In learning-based image compression, Ballé et al.[3] pioneered the use of an encoder-decoder architecture with entropy coding for latent feature representation, enabling model optimization. Progress by Ballé et al.[4] reduced latent feature redundancy through adaptive means and variances. Cheng et al.[7] improved accuracy using Gaussian Mixture Models (GMM) with hyper-priors for GMM parameters. Minnen et al.[5, 12, 32, 33] furthered this with auto-regressive models in entropy coding, decreasing redundancy. However, the sequential nature of auto-regressive methods limited parallel processing, prompting developments like grouped Hyper channel-wise[31] and checkerboard[17] auto-regressive models to enhance speed without losing compression efficiency.

3 METHODOLOGY

3.1 Problem Formulation

The Camera Imaging Pipeline converts a RAW image into a compressed RGB format through a structured process. This process involves three primary stages: conversion from RAW to RGB, compression of the RGB image into a bitstream, and decompression of the bitstream to an RGB image. The aim is to optimize storage and transmission efficiency while preserving image quality. Formally, the pipeline is described by the following sequence of operations:

$$\begin{aligned} \mathbf{r} &= \mathcal{R}(\mathbf{x}) \\ \mathbf{b} &= \mathcal{C}(\mathbf{r}) \\ \mathbf{o} &= \mathcal{D}(\mathbf{b}) \end{aligned} \quad (1)$$

In this formulation, \mathbf{x} is the input RAW image, \mathcal{R} denotes the RAW-to-RGB conversion, \mathbf{r} represents RGB image, \mathcal{C} represents the compression into a bitstream, \mathcal{D} signifies the decompression to RGB, \mathbf{b} is bitstream, and \mathbf{o} is the final RGB output. The end-to-end pipeline encapsulates the entirety of the imaging process, from initial capture to final output, in a single integrated model.

3.2 Architecture of RealCamNet

To enhance the performance of camera imaging pipelines, we propose *RealCamNet*, a novel end-to-end pipeline designed for real-world camera imaging. This pipeline integrates both encoder and decoder components to facilitate effective image processing. The operational framework of *RealCamNet* is formalized as follows:

$$\begin{aligned} \mathbf{y} &= \mathcal{E}(\mathbf{x}; \boldsymbol{\phi}) \\ \hat{\mathbf{y}} &= \mathcal{Q}(\mathbf{y}) \\ \mathbf{o} &= \mathcal{D}(\hat{\mathbf{y}}; \boldsymbol{\theta}) \end{aligned} \quad (2)$$

where \mathbf{x} denotes the input RAW image, and \mathbf{o} represents the output reconstructed RGB image, $\boldsymbol{\phi}$ represents the weights of the encoder \mathcal{E} , while $\boldsymbol{\theta}$ denotes the weights of the decoder \mathcal{D} .

The encoder \mathcal{E} in our pipeline is tasked with performing RAW to RGB conversion and RGB compression encoding. It comprises two principal modules:

- (1) The Coordinate-Aware Distortion Restoration (CADR) module, is aimed at correcting coordinate-dependent distortions such as vignetting, noise, and dark shading, which result from optical and manufacturing defects. The term 'Coordinate' refers to the pixel coordinates in the original RAW image.
- (2) The Coordinate-Independent Mapping Compression (CIMC) module, is responsible for executing spatially invariant operations including global and local tone mapping, denoising, demosaicing, and image feature compression.

For the encoding process, the RAW image \mathbf{x}_{raw} undergoes UnpixelShuffle to achieve channel stacking, converting the spatial dimension RGGB to the channel dimension RGGB image \mathbf{x}_s . Subsequently, \mathbf{x}_s and the coordinate code \mathbf{x}_c are fed into the CADR module to perform position-related restoration, thereby addressing distortions like vignetting and dark shading. A Color Prior Extraction (CPE) module is then utilized to extract color prior information, aiding the CIMC module in performing the tone mapping and compression processes from RAW to RGB. A series of CIMC modules

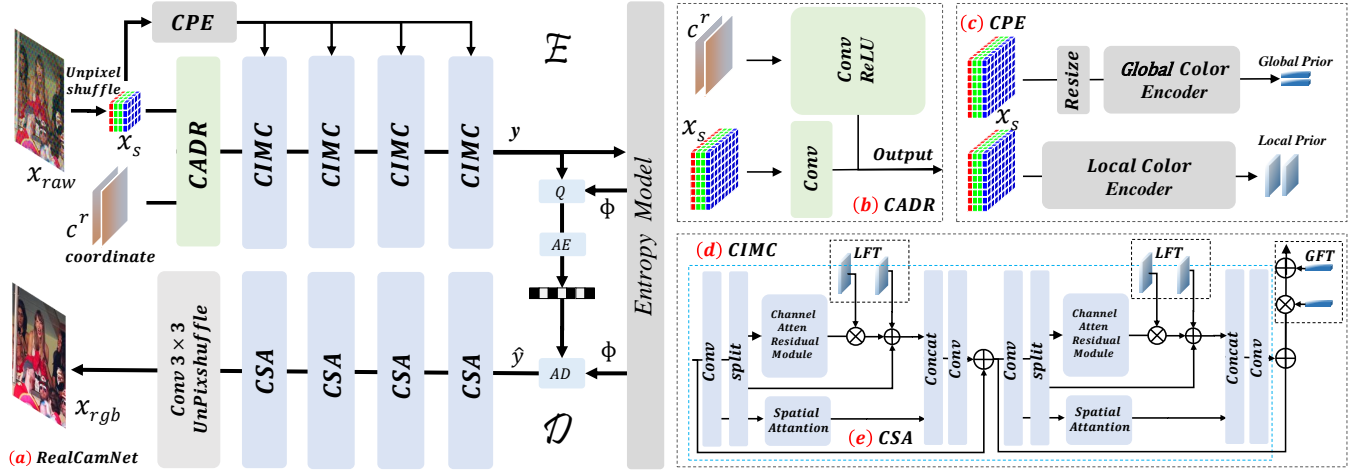


Figure 3: Ours Framework. The encoder E of RealCamNet proposes CADR to restore coordinate-related distortion and builds CIMC to complete coordinate-independent functions (such as global and local tone mapping, denoising, and feature compression). The decoder D of RealCamNet proposes CSA to decode the decoded features and restore the RGB image. LFT is local feature transform, and GFT is global feature transform.

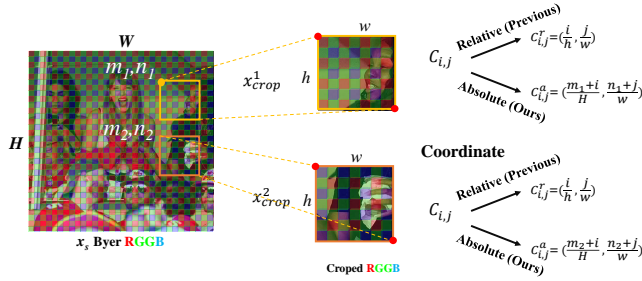


Figure 4: Compared with previous methods that can only encode the relative coordinates of the cropped image, our method calculates the absolute coordinates of the cropped RAW image. Therefore our method can recover fixed position-type distortion in the image.

are employed for image feature compression and tone mapping, followed by the encoding of tone-mapped and compressed latent features into a bitstream via the Entropy Model.

In the decoding phase, the Entropy Model first converts the bitstream back into latent features. These features are processed through the concatenated Channel Spatial Attention (CSA) module, which recovers detailed image information from the compact latent features. Finally, the UnpixelShuffle module is used to reconstruct the RGB image from these features.

This full framework is shown in Fig. 3, illustrates the detail of RealCamNet, which is structured into three main components: encoder, entropy model, and decoder, all of which are jointly trained in an end-to-end manner.

3.3 Coordinate-Aware Distortion Restoration

We present the detailed architecture of the proposed Coordinate-Aware Distortion Restoration (CADR) module, depicted in Fig. 3 (b). The CADR module achieves effective coordinate awareness by integrating the current pixel coordinates, thus facilitating the learning of coordinate-dependent distortion restoration.

Initially, we introduce the previous coordinate calculate method [49]. Since it is impractical to train the network using full-size RAW x_s in $\mathbb{R}^{H,W,4}$ images, typically exceeding 4000×6000 dimensions. Previous methods pre-crop the RAWs into a small cropped RAW dataset, represented as $x_{crop} \in \mathbb{R}^{h,w,4}$, which is suitable for training neural networks. For the cropped image $x_{crop} \in \mathbb{R}^{h,w,4}$, this method ascertains the pixel relative coordinates for each point and stores this coordinate map $c^r \in \mathbb{R}^{h,w,2}$. The calculation method is shown in Eq.3:

$$c_{i,j}^r = \left(\frac{i}{h}, \frac{j}{w} \right) \quad (3)$$

As illustrated in Fig.4, $c_{i,j}^r$ details are presented. Consider that x_{crop}^1 and x_{crop}^2 are randomly cropped from the same RAW image. For left-top coordinate (m_1, n_1) in x_{crop}^1 and left-top coordinate (m_2, n_2) in x_{crop}^2 , the relative coordinates are consistently $c_r^{0,0} = (0, 0)$. However, the absolute coordinates are distinct, with $c_a^{0,0} = \left(\frac{m_1}{H}, \frac{n_1}{W} \right)$ for the first and $c_a^{0,0} = \left(\frac{m_2}{H}, \frac{n_2}{W} \right)$ for the second. Clearly, c_r fails to capture the true coordinates, whereas the c_a successfully retains the actual positional information. This straightforward design facilitates effective distortion restoration.

To address this limitation, we propose a simple method of coordinate embedding to capture the genuine global coordinates of the RAW image. This entails modifying the coordinate computation approach. We first ascertain the coordinates of the top-left pixel of x_{crop} relative to x_s , denoted as $\langle m, n \rangle$. Next, calculate the absolute coordinates $c_{i,j}^a, c_{i,j}^a$ records the coordinate position of each




Figure 5: The detail of the Color Prior Encoder (CPE).

pixel relative to the original RAW, $c_{i,j}^a$. This method of coordinate embedding allows our CADR to accurately perceive the pixel's spatial location, thereby effectively restoring coordinate-related distortions.

$$c_{i,j}^a = \left(\frac{i+m}{H}, \frac{j+n}{W} \right) \quad (4)$$

Upon acquiring the absolute coordinate $c_{i,j}^a$, we integrate the coordinate data into the encoder to achieve coordinate-aware distortion restoration. Initially, the stacked RGGB image x_{crop} is input into a 3×3 convolution layer to extract latent feature x_h . Concurrently, the coordinate information $u_{i,j}^n, v_{i,j}^n$ is processed through a Convolution-ReLU sequence to derive the potential coordinate embeddings x_e . To facilitate coordinate-aware enhancement, the potential coordinate embedding x_e is multiplied by x_s to obtain enhanced features x_o . The entire process is encapsulated in Eq.5.

$$\begin{aligned} x_h &= \text{Conv}(x_s) \\ x_e &= \text{ReLU}(\text{Conv}(c^a)) \\ x_o &= x_h \cdot x_e \end{aligned} \quad (5)$$

3.4 Color Prior Encoder

Figure 5 show the two-fold structure of our Color Prior Extraction (CPE) module: Global prior extraction and Local prior extraction. The global prior extraction process involves downsampling the original 4000×4000 RAW image to a 256×256 representation, preserving global features while reducing computational load, as defined by:

$$G_p = \mathcal{F}_g(\text{downsample}(x_s)), \quad (6)$$

where x_s is the original RAW image, \mathcal{F}_g represents the global color encoder, and G_p denotes the global color prior.

In parallel, local prior extraction focuses on the details within a specific cropped region by applying the local color encoder to the cropped 256×256 RGGB image, described by:

$$L_p = \mathcal{F}_l(x_{crop}), \quad (7)$$

where x_{crop} symbolizes the cropped RAW image, \mathcal{F}_l is the local color encoder, and L_p represents the local color prior. The details of \mathcal{F}_l and \mathcal{F}_g are introduced in the appendix.

3.5 Channel-Spatial Attention

The Channel-Spatial Attention (CSA) mechanism forms an integral part of the Coordinate-independent Mapping Compression (CIMC) module. Inspired by the principles of intra-frame coding

in traditional video compression, the CSA is designed to identify and leverage non-local redundancy within feature representations, promoting more efficient compression. Within the encoder, CSA utilizes an attention mechanism to aggregate information across both spatial and channel dimensions, allowing for the reduction of superfluous data and the distillation of features into compact representations that are conducive to improved compression efficacy. During the decoding phase, CSA plays a crucial role in reconstructing non-local reference features from the compressed feature set. This functionality is vital for restoring the quality of the reconstructed image and ensuring the fidelity of the output relative to the original input.

3.5.1 Detail of CSA. We delve into the computational intricacies of the CSA (Channel Spatial Attention) as follows. The initial input feature x_{in} undergoes a transformation via a convolution layer:

$$x' = \text{Conv}(x_{in}) \quad (8)$$

Subsequently, the processed feature x' is bifurcated into two distinct components x_1 and x_2 :

$$x_1, x_2 = \text{Split}(x') \quad (9)$$

The first component, x_1 , is subjected to the Channel-Wise Residual Attention (CWRA) yielding x_3 :

$$x_{ca} = \text{CWRA}(x_1) \quad (10)$$

Simultaneously, the second component, x_2 , traverses through the Spatial-Wise Attention (SWA) module, producing x_{sa} :

$$x_{sa} = \text{SWA}(x_2) \quad (11)$$

In our spatial-wise attention module, we employ the Swin Transformer. The fusion of x_{sa} and x_{ca} is accomplished via concatenation, followed by a convolutional layer to get output feature x_o :

$$x_o = \text{Conv}(\text{Concat}(x_{sa}, x_{ca})) \quad (12)$$

This approach synergizes channel and spatial attention, enhancing feature refinement and facilitating the transformation of RAW image features into a compact RGB latent representation. The result is an improved reconstruction quality with reduced redundancy, essential for efficient image processing tasks.

3.6 Coordinate-Independent Mapping Compression

3.6.1 Overview of CIMC. The Coordinate-independent Mapping Compression (CIMC) module is designed to compress RAW image features into compact RGB latent representations. Employing attention mechanisms and feature transformations, CIMC aims to enhance reconstruction quality and reduce data redundancy.

CIMC consists of three key components: the Channel-Spatial Attention (CSA) module, the Local Feature Transformation (LFT) module, and the Global Feature Transformation (GFT) module. LFT and GFT can use the color prior extracted by the CPE module for tone mapping. CSA promotes effective compression of latent features, as shown in Fig.3 (d).

3.6.2 *Details of CIMC.* Incorporating both *GFT* and *LFT* within the Encoder, *CIMC* extends *CSA* to perform global and local tone mapping essential for the camera imaging pipeline. Positioned after the *CSA* channel attention stage, *LFT* serves to refine features (Fig. 3):

$$x_{ca} = LFT(x_{ca}, L_p) \quad (13)$$

Subsequent iterations of *CSA* and *LFT* result in the feature x_c :

$$x_c = LFT(CSA(LFT(CSA(x_h), L_p)), L_p) \quad (14)$$

The concluding phase entails the transformation of x_c by *GFT* to yield the output x_o , representing the compressed RGB latent features:

$$x_o = GFT(x_c, G_p) \quad (15)$$

where G_p, L_p is the output of the CPE module. The calculation of *GFT* and *LFT* are both $y = \alpha x + \beta$. The difference is that α and β are globally or locally adaptive. This module allows *CIMC* to enhance color fidelity through integrated local and global feature transformations.

4 EXPERIMENT

4.1 Implementation Details

To train RealCamNet, we created a comprehensive RAW-RGB dataset. Our dataset, with 4507 RAW-RGB image pairs at 6000×4000 resolution, covers diverse scenes including animals, landscapes, and architecture. For evaluation, we employ 450 pairs, with the remainder for training. Further dataset specifics are in the Appendix. For optimization, we employed the Adam optimizer [25, 43] with an initial learning rate of 1×10^{-4} , implementing a multi-step rate reduction strategy. The RD-formula loss, as per [7], was used, testing λ values within [0.1, 0.025, 0.01, 0.005] to accommodate different bitrates. The training was conducted on an i9-13900K CPU and an RTX 4090 GPU, with batch size set at 8.

4.2 Metrics

Evaluation metrics used in our study include PSNR (Peak Signal-to-Noise Ratio) [38], MS-SSIM (Multi-Scale Structural Similarity Index) [42], ΔE (a measure of color difference) [8], and LPIPS (Learned Perceptual Image Patch Similarity) [47]. PSNR measures image reconstruction quality, MS-SSIM evaluates image fidelity across multiple scales, ΔE assesses color accuracy, and LPIPS quantifies perceptual similarity using deep learning models. For clarity, MS-SSIM was converted to $-10 \log_{10}(1 - \text{MS-SSIM})$. The performance across metrics was gauged using Bjøntegaard Delta (BD) metrics, encompassing BD-PSNR, BD-MSSSIM, BD- ΔE , and BD-LPIPS.

4.3 Rate-Distortion Performance

To assess the efficacy of our newly proposed RealCamNet, we juxtapose its performance with the leading multi-stage separation schemes. The comparative analysis encompasses two principal pipelines: (1) a learning-based ISPNet combined with the Universal Video Coding (VVC) framework, and (2) ISPNet integrated with a learning-based image compression approach. Fig.6 delineates the rate-distortion outcomes across our test dataset. A thorough evaluation is conducted employing a suite of metrics—peak signal-to-noise ratio (PSNR), multi-scale structural similarity index (MS-SSIM),

color fidelity (ΔE), and learning-based perceptual image patch similarity (LPIPS)—to gauge both the objective and perceptual quality, thus offering a holistic assessment of the varied pipelines.

With PyNet+VTM set as the anchor, we delve into the performance nuances of each pipeline using diverse metrics, calculating the Bjøntegaard Delta (BD) metrics across the metrics-BPP curve, inclusive of BD-PSNR, BD-MSSSIM, BD- ΔE , and BD-LPIPS. The empirical data reveal that our RealCamNet, at an equivalent bitrate, transcends the conventional SOTA multi-stage, learning-based pipelines, marking improvements of 2.26dB in PSNR, 0.71dB in MS-SSIM, 0.01 in LPIPS, and 0.9187 in ΔE . Additionally, when pitted against the learning-based ISPNet in conjunction with the Cascaded Universal Video Coding (VVC) framework, our method consistently exhibits superior performance in metrics like PSNR at comparable bitrates, underscoring the robustness and stability of our proposed approach.

To elucidate the performance of RealCamNet, we computed the BD-Rate based on the rate-distortion curve, revealing a significant enhancement over the best-existing pipeline, with a BD-Rate improvement of 9.53%. These findings, presented in Table 1, underscore the comprehensive performance superiority of our approach on the test set.

4.4 Ablation Studies of CIMC

In Fig.9, we assess the efficacy of the **CIMC** module, which comprises *CSA*, *LFT*, and *GFT* components. The **CIMC** module enhances tone mapping through novel *LFT* and *GFT* while optimizing feature compression with *CSA*. Quantitative analysis of the **CIMC** and its individual *CSA*, *LFT*, and *GFT* components is presented in Table 2.

4.5 Ablation Studies of CADR

To demonstrate the effectiveness of our proposed Coordinate-Aware Distortion Restoration module **CADR**, we compared the model without the **CADR** module with the model adding the **CADR** module. The results are shown in Fig.9. The addition of our **CADR** module brought huge improvements. We show the quantitative gain on BD-PSNR in Table 2.

4.6 Complexity and Qualitative Results.

We compare the computational complexity of different methods. As shown in Table6 and Figure8, our method can decode images with a resolution of 1024×1024 at an inference speed of 16.8fps. Our approach not only achieves superior performance but also demonstrates lower computational complexity and faster inference speed. This dual advantage underscores the method's efficiency, marrying high-quality results with expeditious processing.

4.7 Visual Results

Fig. 7 presents a visual comparison of the reconstructed images using our method against Pipeline-1 (ISPNet coupled with the classic VVC standard, VTM 12.1) and Pipeline-2 (ISPNet integrated with a learning-based compression network). Our method exhibits superior performance in preserving intricate textures, as evidenced by clearer feather outlines, and achieves enhanced color fidelity.

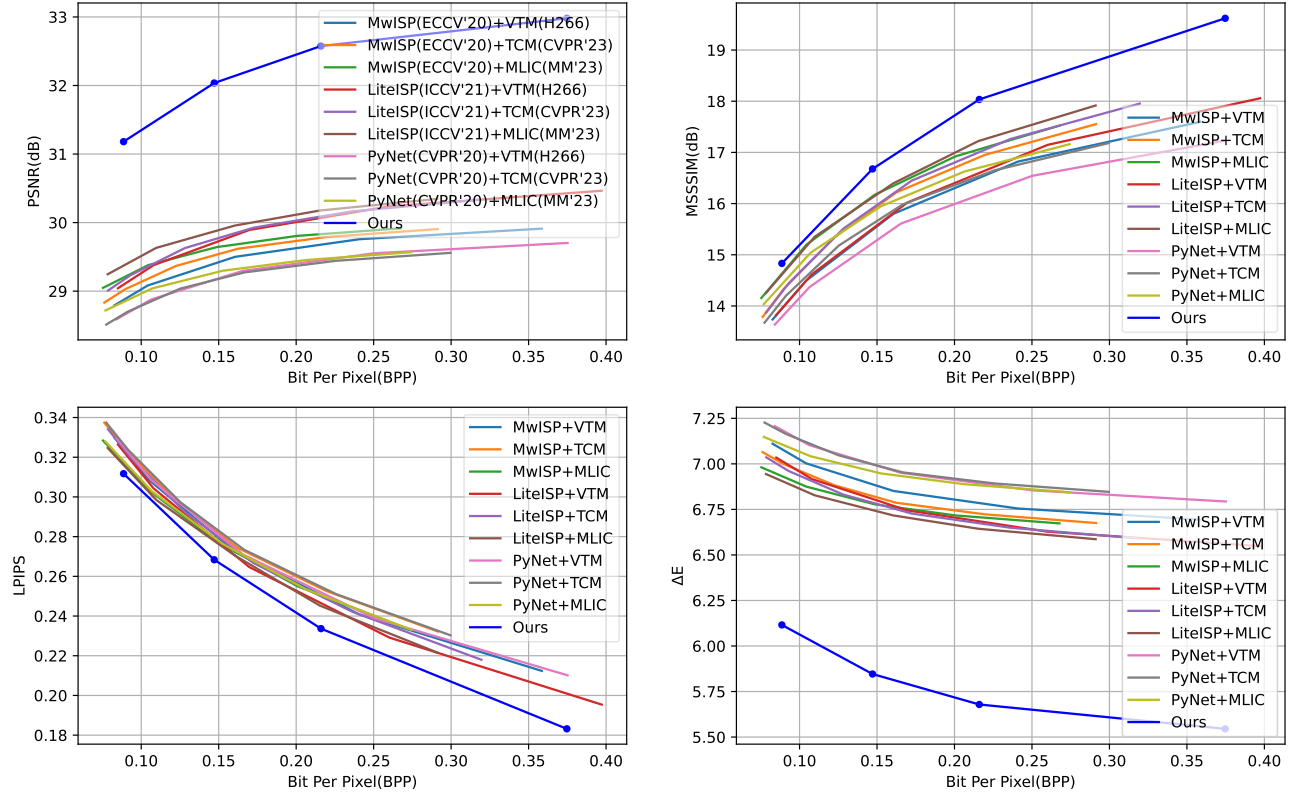


Figure 6: Quantitative Rate-Distortion curve results.

Table 1: Quantitative results. We compare with state-of-the-art ISPNet and image compression methods, including learning-based methods: PyNet(CVPR'20)[19], LiteISPNet(ICC'21)[48], MwISPNet(ECCV'20)[18], MLIC(ACMMM'23)[22], TCM(CVPR'23)[31] and the most advanced traditional image compression method VTM/H.266. We show BD-Rate, BD-PSNR, BD-MSSSIM, BD- ΔE and BD-LPPHS for all methods. We use PyNet+VTM as anchor.

Method	Params (M) \downarrow	FLOPs (G) \downarrow	Enc Time (s) \downarrow	Dec Time (s) \downarrow	BD - Rate \downarrow	BD- PSNR(db) \uparrow	BD- MSSSIM(db) \uparrow	BD- LPIPS \downarrow	BD- ΔE \downarrow
PyNet+VTM[1, 19]	-	-	196.94	0.1433	0.0000	0.0000	0.0000	0.0000	0.0000
PyNet+TCM[19, 31]	92.51	2354	0.1582	0.1405	-12.1248	-0.0286	0.3270	0.0031	0.0079
PyNet+MLIC[19, 22]	164.03	3108	0.1695	0.2155	-20.2375	0.0776	0.5603	-0.0025	-0.0314
MwISP+VTM[1, 18]	-	-	196.93	0.1433	-10.2986	0.2262	0.2868	-0.0012	-0.1071
MwISP+TCM[18, 31]	74.18	1026	0.1615	0.1405	-20.7415	0.3292	0.6163	0.0022	-0.1658
MwISP+MLIC[18, 22]	145.7	1780	0.1727	0.2155	-28.7059	0.4408	0.8593	-0.0032	-0.2076
LiteISP+VTM[1, 48]	-	-	196.93	0.1433	-12.5324	0.5880	0.3900	-0.0052	-0.2022
LiteISP+TCM[31, 48]	53.98	940	0.1522	0.1405	-22.0966	0.5911	0.7029	-0.0010	-0.2094
LiteISP+MLIC[22, 48]	125.5	1694	0.1635	0.2155	-29.5519	0.7079	0.9223	-0.0057	-0.2513
Ours	49.01	357	0.0703	0.0592	-39.0842	2.9603	1.6392	-0.0162	-1.1709

4.8 Receptive Field Analysis

We analyze the effectiveness of each module by calculating the effective receptive field. The analysis of receptive fields for different inputs indicates that the model's output for Global downsamples RAW encompasses a global receptive field(Fig.10 (b)), thus facilitating the extraction of a global color before the entire image. Furthermore, the output for Cropped RAW exhibits a non-local

receptive field(Fig.10 (c)), demonstrating that CSA (Channel Spatial Attention) is capable of harnessing non-local information to enhance the compactness of latent. In contrast, coordinate embedding is local to the input receptive field(Fig.10 (d)) because perceiving the coordinate information of the current position enables high-quality position-dependent distortion restoration.

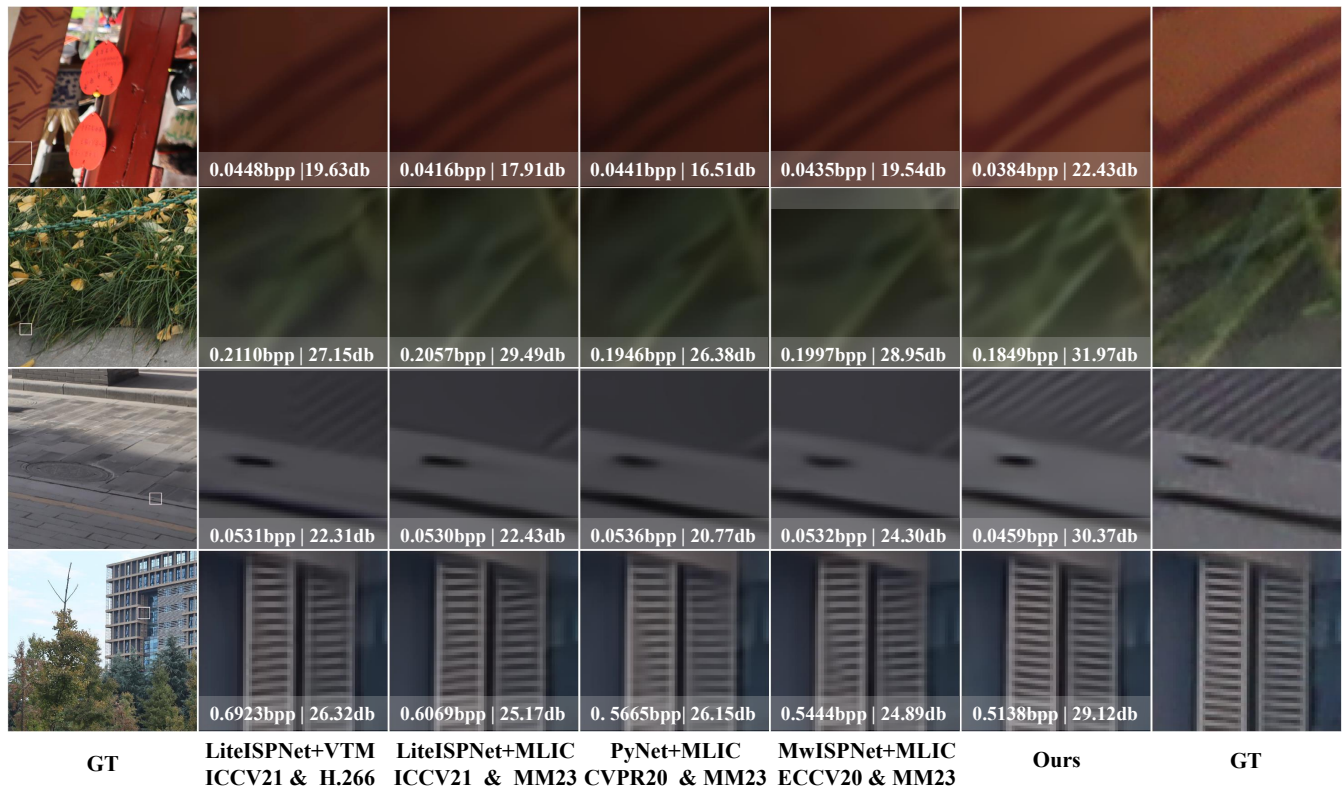


Figure 7: Compared with the state-of-the-art methods, our framework, optimized end-to-end, demonstrates significant enhancements in imaging systems’ performance. It offers comprehensive benefits, encompassing reduced bit rate, augmented color fidelity, and elevated PSNR.




Figure 8: Results of performance and inference time.




Figure 9: Ablation in Rate-Distortion curve.

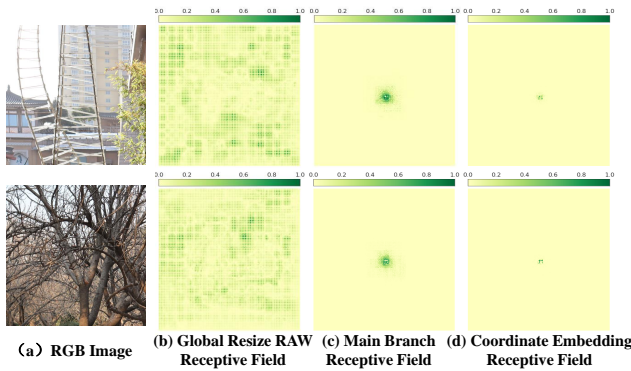
5 CONCLUSION

This study unveils an innovative end-to-end real-world camera imaging pipeline, surpassing existing methods in both performance and efficiency and pioneering a new benchmark for imaging pipeline design. Through a detailed analysis of the ISP and compression

workflows, we introduced a coordinate-independent mapping compression module aimed at optimizing feature compression and tone mapping, alongside a coordinate-aware restore module dedicated to restoring coordinate-specific distortions. The bespoke dataset,

Table 2: Quantitative ablation results. Results are presented with 'Base' serving as the benchmark anchor.

Base	CADR	CSA	GFT	LFT	Params (M)	FLOPs (G)	BD-PSNR(db)
✓					48.511	79.810	0.0000
✓	✓				48.561	83.048	0.4946
✓	✓	✓			46.071	71.059	0.7475
✓	✓	✓	✓		47.290	77.438	2.6545
✓	✓	✓	✓	✓	49.010	89.176	3.1716

**Figure 10: Receptive field analysis. Our global branch can effectively extract global information to assist in more accurate color restoration, while coordinate awareness requires only a local receptive field to achieve recovery of coordinate-related distortion types.**

crafted for camera imaging pipeline evaluation, sets a robust benchmark, catalyzing future research. Our exhaustive evaluation not only confirms the pipeline’s efficacy but also underscores its potential to revolutionize future digital imaging technologies.

REFERENCES

- [1] 2020. *Versatile Video Coding (H.266/VVC)*. Technical Report. Joint Video Experts Team (JVET). ITU-T Recommendation H.266 and ISO/IEC 23090-3.
- [2] Theo Adrai, Guy Ohayon, Michael Elad, and Tomer Michaeli. 2024. Deep optimal transport: A practical algorithm for photo-realistic image restoration. *Advances in Neural Information Processing Systems* 36 (2024).
- [3] Johannes Ballé, Valero Laparra, and Eero P Simoncelli. 2016. End-to-end optimized image compression. *arXiv preprint arXiv:1611.01704* (2016).
- [4] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. 2018. Variational image compression with a scale hyperprior. *arXiv preprint arXiv:1802.01436* (2018).
- [5] Jean Bégaint, Fabien Racapé, Simon Feltman, and Akshay Pushparaja. 2020. Compressai: a pytorch library and evaluation platform for end-to-end compression research. *arXiv preprint arXiv:2011.03029* (2020).
- [6] Antoni Buades, Bartomeu Coll, and Jean-Michel Morel. 2011. Non-local means denoising. *Image Processing On Line* 1 (2011), 208–212.
- [7] Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. 2020. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 7939–7948.
- [8] Commission Internationale de l’Eclairage (CIE). 2004. Colorimetry. CIE 15:2004. Technical Report, 3rd edition.
- [9] Jonathan W Cunningham, Pulkit Singh, Christopher Reeder, Brian Claggett, Pablo M Marti-Castellote, Emily S Lau, Shaan Khurshid, Puneet Batra, Steven A Lubitz, Mahnaz Maddah, et al. 2024. Natural language processing for adjudication of heart failure in a multicenter clinical trial: a secondary analysis of a randomized clinical trial. *JAMA cardiology* 9, 2 (2024), 174–181.
- [10] Linhui Dai, Xiaohong Liu, Chengqi Li, and Jun Chen. 2020. Awnet: Attentive wavelet network for image isp. In *Computer Vision—ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*. Springer, 185–201.
- [11] Paul Debevec and Simon Gibson. 2002. A tone mapping algorithm for high contrast images. In *13th eurographics workshop on rendering: Pisa, Italy. Citeseer*.
- [12] Emilien Dupont, Hrushikesh Loya, Milad Alizadeh, Adam Goliński, Yee Whye Teh, and Arnaud Doucet. 2022. Coin++: Neural compression across modalities. *arXiv preprint arXiv:2201.12904* (2022).
- [13] Fredo Durand and Julie Dorsey. 2000. Interactive tone mapping. In *Rendering Techniques 2000: Proceedings of the Eurographics Workshop in Brno, Czech Republic, June 26–28, 2000 11*. Springer, 219–230.
- [14] Burhan Ergen. 2012. *Signal and image denoising using wavelet transform*. InTech London, UK.
- [15] Hansen Feng, Lizhi Wang, Yuzhi Wang, and Hua Huang. 2022. Learnability Enhancement for Low-light Raw Denoising: Where Paired Real Data Meets Noise Modeling. In *Proceedings of the 30th ACM International Conference on Multimedia (<conf-loc>, <city>Lisboa</city>, <country>Portugal</country>, </conf-loc>)* (MM ’22). Association for Computing Machinery, New York, NY, USA, 1436–1444. <https://doi.org/10.1145/3503161.3548186>
- [16] Hansen Feng, Lizhi Wang, Yuzhi Wang, and Hua Huang. 2022. Learnability enhancement for low-light raw denoising: Where paired real data meets noise modeling. In *Proceedings of the 30th ACM International Conference on Multimedia*. 1436–1444.
- [17] Dailan He, Yaoyan Zheng, Baocheng Sun, Yan Wang, and Hongwei Qin. 2021. Checkerboard context model for efficient learned image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 14771–14780.
- [18] Andrey Ignatov, Radu Timofte, Zhilu Zhang, Ming Liu, Haolin Wang, Wangmeng Zuo, Jiawei Zhang, Ruimao Zhang, Zhanglin Peng, Sijie Ren, et al. 2020. 2020 challenge on learned image signal processing pipeline. In *Computer Vision—ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*. Springer, 152–170.
- [19] Andrey Ignatov, Luc Van Gool, and Radu Timofte. 2020. Replacing mobile camera isp with a single deep learning model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. 536–537.
- [20] Andrey Ignatov, Luc Van Gool, and Radu Timofte. 2020. Replacing mobile camera isp with a single deep learning model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. 536–537.
- [21] Wooseok Jeong and Seung-Won Jung. 2022. RAWtoBit: A Fully End-to-end Camera ISP Network. In *European Conference on Computer Vision*. Springer, 497–513.
- [22] Wei Jiang, Jiayu Yang, Yongqi Zhai, Peirong Ning, Feng Gao, and Ronggang Wang. 2023. Mlic: Multi-reference entropy model for learned image compression. In *Proceedings of the 31st ACM International Conference on Multimedia*. 7618–7627.
- [23] Xin Jin, Jia-Wen Xiao, Ling-Hao Han, Chunle Guo, Ruixun Zhang, Xialei Liu, and Chongyi Li. 2023. Lighting every darkness in two pairs: A calibration-free pipeline for raw denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 13275–13284.
- [24] Aditya Joshi, Raj Dabre, Diptesh Kanojia, Zhuang Li, Haolan Zhan, Gholamreza Haffari, and Doris Dippold. 2024. Natural language processing for dialects of a language: A survey. *arXiv preprint arXiv:2401.05632* (2024).
- [25] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [26] Ann-Christin Klemenz, Lasse Albrecht, Mathias Manzke, Antonia Dalmer, Benjamin Böttcher, Alexey Surov, Marc-André Weber, and Felix G Meinel. 2024. Improved image quality in CT pulmonary angiography using deep learning-based image reconstruction. *Scientific Reports* 14, 1 (2024), 2494.
- [27] Rakesh Kumar, Pooja Kumbharkar, Sandeep Vanam, and Sanjeev Sharma. 2024. Medical images classification using deep learning: a survey. *Multimedia Tools and Applications* 83, 7 (2024), 19683–19728.
- [28] Marc Lebrun. 2012. An analysis and implementation of the BM3D image denoising method. *Image Processing On Line* 2 (2012), 175–213.
- [29] Xin Li, Bahadır Gunturk, and Lei Zhang. 2008. Image demosaicing: A systematic survey. In *Visual Communications and Image Processing 2008*, Vol. 6822. SPIE, 489–503.
- [30] Zhetong Liang, Jianrui Cai, Zisheng Cao, and Lei Zhang. 2021. Cameranet: A two-stage framework for effective camera isp learning. *IEEE Transactions on Image Processing* 30 (2021), 2248–2262.
- [31] Jinming Liu, Heming Sun, and Jiro Katto. 2023. Learned Image Compression with Mixed Transformer-CNN Architectures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1–10.
- [32] Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. 2018. Conditional probability models for deep image compression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4394–4402.
- [33] David Minnen, Johannes Ballé, and George D Toderici. 2018. Joint autoregressive and hierarchical priors for learned image compression. *Advances in neural*

- information processing systems* 31 (2018).
- [34] Behnoosh Mohammadzadeh, Jules Françoise, Michèle Gouiffès, and Baptiste Caramiaux. 2024. Studying Collaborative Interactive Machine Teaching in Image Classification. In *Proceedings of the 29th International Conference on Intelligent User Interfaces*. 195–208.
- [35] Shiyu Ning, Hongteng Xu, Li Song, Rong Xie, and Wenjun Zhang. 2018. Learning an inverse tone mapping network with a generative adversarial regularizer. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1383–1387.
- [36] DN Kiran Pandiri, R Murugan, and Tripti Goel. 2024. Smart soil image classification system using lightweight convolutional neural network. *Expert Systems with Applications* 238 (2024), 122185.
- [37] Priyanshu Priya, Mauajama Firdaus, and Asif Ekbal. 2024. Computational Politeness in Natural Language Processing: A Survey. *Comput. Surveys* (2024).
- [38] David Salomon. 2004. *Data Compression: The Complete Reference* (4 ed.). Springer. Chapter on Image Compression, specifically the section on PSNR.
- [39] Alireza Shamshiri, Kyeong Rok Ryu, and June Young Park. 2024. Text mining and natural language processing in construction. *Automation in Construction* 158 (2024), 105200.
- [40] Ling Shao and Amin Ur Rehman. 2014. Image demosaicing using content and colour-correlation analysis. *Signal processing* 103 (2014), 84–91.
- [41] Anthony A Tanbakuchi, Arjen van der Sijde, Bart Dillen, Albert JP Theuwissen, and Wim de Haan. 2003. Adaptive pixel defect correction. In *Sensors and Camera Systems for Scientific, Industrial, and Digital Photography Applications IV*, Vol. 5017. SPIE, 360–370.
- [42] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. 2003. Multiscale structural similarity for image quality assessment. *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003* 2 (2003), 1398–1402.
- [43] Blake Woodworth, Kumar Kshitij Patel, Sebastian Stich, Zhen Dai, Brian Bullins, Brendan McMahan, Ohad Shamir, and Nathan Srebro. 2020. Is local SGD better than minibatch SGD?. In *International Conference on Machine Learning*. PMLR, 10334–10343.
- [44] Yifan Xu, Mengdan Zhang, Chaoyou Fu, Peixian Chen, Xiaoshan Yang, Ke Li, and Changsheng Xu. 2024. Multi-modal queried object detection in the wild. *Advances in Neural Information Processing Systems* 36 (2024).
- [45] Sai Harsha Yelleni, Deepshikha Kumari, PK Srijith, et al. 2024. Monte Carlo DropBlock for modeling uncertainty in object detection. *Pattern Recognition* 146 (2024), 110003.
- [46] Quan Zhang, Xiaoyu Liu, Wei Li, Hanting Chen, Junchao Liu, Jie Hu, Zhiwei Xiong, Chun Yuan, and Yunhe Wang. 2024. Distilling semantic priors from sam to efficient image restoration models. *arXiv preprint arXiv:2403.16368* (2024).
- [47] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 586–595.
- [48] Zhilu Zhang, Haolin Wang, Ming Liu, Ruohao Wang, Jiawei Zhang, and Wangmeng Zuo. 2021. Learning raw-to-srgb mappings with inaccurately aligned supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 4348–4358.
- [49] Zhilu Zhang, Haolin Wang, Ming Liu, Ruohao Wang, Jiawei Zhang, and Wangmeng Zuo. 2021. Learning raw-to-srgb mappings with inaccurately aligned supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 4348–4358.