

Annual Review of Vision Science

Mobile Computational Photography: A Tour

Mauricio Delbracio,¹ Damien Kelly,¹
Michael S. Brown,² and Peyman Milanfar¹

¹Google Research, Mountain View, California 94043, USA; email: mdelbra@google.com, damienkelly@google.com, milanfar@google.com

²EECS Department, York University, Toronto, Ontario M6E 3N1, Canada; email: mbrown@eecs.yorku.ca

Annu. Rev. Vis. Sci. 2021. 7:571–604

The *Annual Review of Vision Science* is online at
vision.annualreviews.org

<https://doi.org/10.1146/annurev-vision-093019-115521>

Copyright © 2021 by Annual Reviews.
All rights reserved

ANNUAL
REVIEWS **CONNECT**

www.annualreviews.org

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

Keywords

computational photography, mobile camera, noise reduction, burst processing, super-resolution, low-light imaging, zoom, camera pipeline

Abstract

The first mobile camera phone was sold only 20 years ago, when taking pictures with one's phone was an oddity, and sharing pictures online was unheard of. Today, the smartphone is more camera than phone. How did this happen? This transformation was enabled by advances in computational photography—the science and engineering of making great images from small-form-factor, mobile cameras. Modern algorithmic and computing advances, including machine learning, have changed the rules of photography, bringing to it new modes of capture, postprocessing, storage, and sharing. In this review, we give a brief history of mobile computational photography and describe some of the key technological components, including burst photography, noise reduction, and super-resolution. At each step, we can draw naive parallels to the human visual system.

1. INTRODUCTION AND HISTORICAL OVERVIEW

Modern digital photography has a fascinating and nuanced history (**Figure 1**), punctuated by important advances in both camera sensor technology and algorithms that operate on the captured signals. In this review, we concentrate on the most recent two decades of intense and rapid progress in computational photography. More specifically, we provide an excursion through mobile computational photography, where computational photography has had the largest impact on the daily lives of people across the globe. From news to social media, digital photographs (now overwhelmingly captured on mobile devices) have fundamentally transformed how we capture and remember our world. Indeed, it is not an exaggeration to say that the smartphone (and its camera in particular) has changed the world (see **Figure 2**).

The era of analog (film) photography saw its golden age from the 1930s onward, when giants of the craft, like Ansel Adams, practiced their art of “making a photograph” (Adams 1935) by hand in their custom dark rooms. Remarkably, innovations by Adams and others, such as the dodge and burn technique for high-dynamic-range photography, have persisted in digital and computational photography to this day, albeit in much more formal, algorithmic incarnations. Analog film thus dominated the scene for nearly 50 years but was largely discontinued in 2005, when standalone point-and-shoot digital cameras became dominant and before cellphones had good imaging capabilities.

The first commercial digital cameras appeared in the early 1990s. In 1992, the first digital single-lens reflex (DSLR) cameras entered the market, but they were prohibitively expensive, costing upwards of \$20,000. Unsurprisingly, these early devices failed to capture a large market share. The introduction of complementary metal oxide semiconductor (CMOS) image sensors in 1993 facilitated the development of what became known as camera on a chip. This revolutionary sensor would enable far less expensive devices to be built with proportionally better power efficiency. Yet the technical difficulties in replacing the existing standard of the time (charge-coupled device arrays) were significant. These included noisier pixels and a rolling (rather than global) shutter in

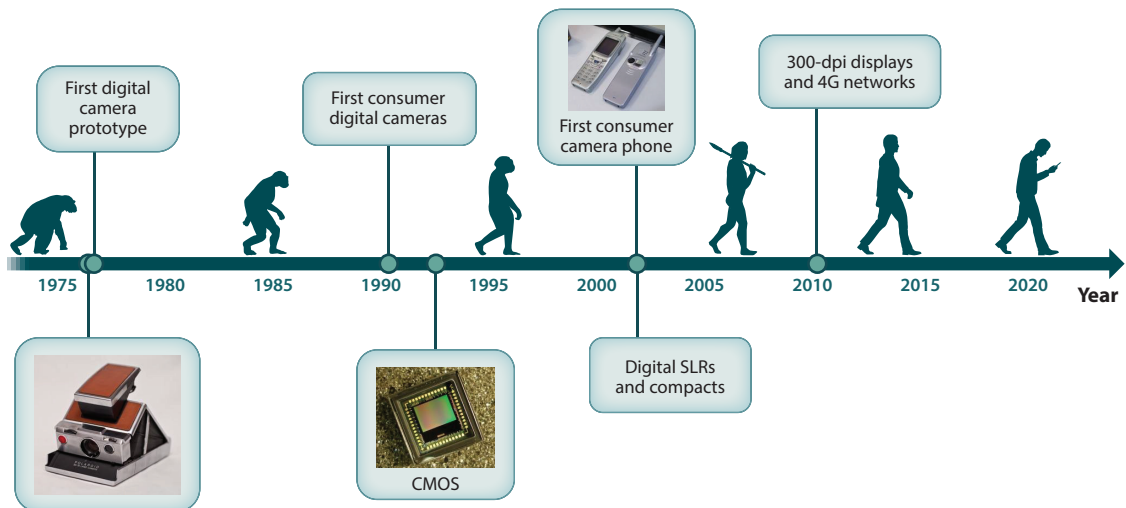


Figure 1

An incomplete timeline of photographic innovations in the past four decades. Abbreviations: CMOS, complementary metal oxide semiconductor; dpi, dots per inch; SLR, single-lens reflex.



Figure 2

The crowds at Saint Peter's Square in Vatican City in 2005 and 2013. Photographs by Luca Bruno/AP, Michael Sohn/AP.

the cheaper CMOS sensors. Because of these technical difficulties, it would take another 10 years before CMOS systems would enable mass production of mobile and digital cameras.

In the early 2000s, the dominant digital devices for photography were single-lens reflex (SLR) and digital point-and-shoot cameras. Today, 20 years later, the prices of these two camera types are roughly in the same proportion as they were then, while the quality of each has significantly improved—meaning that both technologies now present a much better value for the user. In standalone cameras, somewhat surprisingly, this improvement is largely due to better sensors and optics and, to a much lesser degree, to technical advances in software. Even today, standalone cameras still have relatively naive software pipelines. Meanwhile, algorithmic and software innovations have thrived in mobile devices. Why this is the case is an interesting question. While many factors have contributed to this dichotomy—hardware advances on the standalone cameras versus software advances on mobile platforms enabled by significant computing power—one thing remains clear: The requirements imposed by the sleek industrial design and consequent form factors in mobile smartphones have severely limited what can be done with imaging hardware. As a result, the best possible hardware on mobile devices is still often bettered by complementary development and close integration of algorithmic software solutions. In some ways, the situation is not altogether dissimilar to the development of vision in nature. The evolution of the physical form of the eye has had to contend with physical constraints that limit the size, shape, and sensitivity of light gathering in our vision system. In tandem and complementary fashion, the visual cortex has developed the computational machinery to interpolate, interpret, and expand the limits imposed by the physical shape of our eyes—our visual hardware.

The introduction of the iPhone in 2007 was a watershed moment in the evolution of mobile devices and changed the course of both phone and camera technology. Looking back at the early devices, the layperson may conclude that the camera was immediately transformed to become the high-utility application that we know today. This was in fact not the case. What was reinvented for the better at the time were the display and user interface, but not necessarily the camera; indeed, the two-megapixel camera on the first iPhone was far inferior to just about any existing point-and-shoot camera of comparable price at the time.

The year 2010 was pivotal for the mobile camera. A transition to both 4G wireless and 300 dots per inch displays enabled users to finally appreciate their photographs on the screens of their own mobile devices. Indeed, users felt that their phone screens were not only rich enough for the consumption of their own photos, but also sufficient to make photo sharing worthwhile. Meanwhile, significantly faster wireless network speeds meant that individuals could share their photos almost instantaneously (see **Figure 3**). Once viewing and sharing had been improved by leaps and bounds, it became imperative for mobile manufacturers to significantly improve the quality of the captured photos as well. Thus began an emphasis on improved light collection, better dynamic




Figure 3

The year 2010 saw the convergence of two important trends: improved displays and increased wireless speed. These forces conspired to make mobile photography the dominant mode of imaging in the ensuing decade. The speed of 0.01 seconds is expected from 5G, although it has not yet been delivered in practice. Abbreviations: dpi, dots per inch; Mbps, megabits per second.

range, and higher resolution for camera phones so that consumers could use their phones as cameras and communication devices. The added bonus was that users would no longer have to carry both a phone and a camera—they could rely on their smartphone as a multipurpose device.

To achieve this ambitious goal meant overcoming a large number of limitations—challenges that have kept the computational photography community busy for the past decade. In the next section, we begin by highlighting some of the main limitations of the mobile imaging platform as compared to standalone devices, and we go on in the rest of the article to review how these challenges have been met in research and practice.

2. THE MOBILE CAMERA: HARDWARE AND ITS LIMITATIONS

Ideally, a smartphone camera would offer photographic performance on par with a DSLR (or at least a compact point-and-shoot) camera. However, the smartphone camera has several notable disadvantages due to constraints placed on its form factor for it to be integrated into the phone's thin profile. **Figure 4** shows a smartphone with integrated cameras next to a typical DSLR camera.

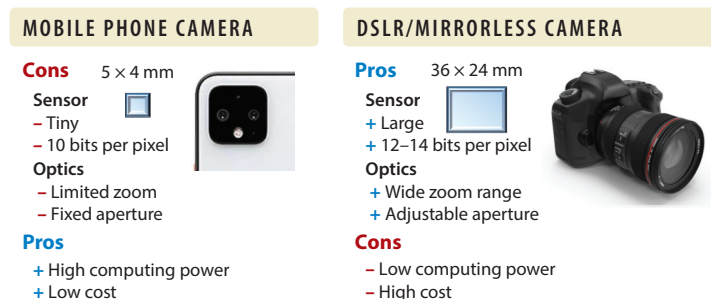


Figure 4

The pros and cons of a smartphone compared to a digital single-lens reflex (DSLR) camera. The most notable differences are the larger sensor and optics available on a DSLR camera. Surprisingly, however, a high-end smartphone has significantly more computing power than most DSLR cameras.

The physical camera sensor and associated lens optics of the smartphone camera are significantly smaller and less flexible than those of a DSLR camera. Yet, while the smartphone's physical hardware is limited, the smartphones have access to much more computing power than is available on a DSLR camera. To draw a rough but stark contrast between the two platforms, a mobile camera's small aperture limits light collection by two orders of magnitude as compared to a typical DSLR camera. Meanwhile, the same mobile device houses roughly two orders of magnitude more computing power. The trade-off of additional computing for more sophisticated imaging hardware is thus inevitable.

In this section, we briefly summarize several key limitations of the mobile phone camera as compared to a DSLR camera.

2.1. Sensor Size and Limited Aperture

The most obvious limitations of a smartphone camera are the size of its sensor and the compactness of its optics. Modern smartphone sensors are roughly 5×4 mm in size, while many DSLR cameras still use full-size 36×24 mm sensors. In addition to a small sensor size, the optics of the mobile phone camera are significantly smaller and less adjustable compared to a typical lens used on a DSLR camera. In addition, most mobile cameras use a compact lens array that has a fixed aperture. The focal length is also limited, leading many phone makers to include two or more cameras with different focal lengths, each serving a different purpose (main, zoom, wide, etc.)

2.2. Noise and Limited Dynamic Range

Image noise can be defined as random unwanted variations of the intensity level of image pixels. In addition to random fluctuations due to thermal agitation in electronics, there exists a permanent, unavoidable source of noise due to the discrete nature of light (photon shot noise).

With its smaller aperture and sensor size, for a given exposure time, a smartphone is able to capture at best a small fraction of the amount of light that would be captured by a DSLR camera. A smaller sensor also means that even less light hits the surface of the sensor when capturing an image. As a result, smartphone cameras often need to apply a nontrivial multiplicative gain to the recorded signal. This gain is controlled by the ISO setting—a higher ISO number implies an increase in the gain factor, which amplifies the sensor noise. Consequently, the smartphone camera images produced at the sensor are markedly more noisy than images captured with DSLR sensors.

Another notable difference between DSLR and smartphone cameras is the dynamic range of the sensor, which is defined as the ratio between the full well capacity of a pixel's photodiode at maximum gain and its noise (read noise). In practice, this defines the brightest and darkest parts of the scene that can be captured without clipping or saturation. The dynamic range is directly correlated to the pixel size. A DSLR pixel's photodiode is roughly 4 microns in width, while a smartphone sensor is closer to 1.5 microns or less. This means that pixels of a smartphone sensor have a much smaller well capacity, and therefore, the maximum amount of electrical current that they can capture at each photodiode is reduced. As a result, whereas a DSLR camera can effectively encode anywhere between 4,096 (12 bits) and 16,384 (14 bits) tones per pixel, a typical smartphone camera sensor is limited to 1,024 (10 bits) tonal values per pixel.

2.3. Limited Depth of Field

The depth of field (DoF) defines the region in the image of the scene where objects appear sharp. The DoF can be controlled by the camera's focal length and aperture. The wider is the aperture,




Figure 5

A typical camera sensor with a color filter array layout (Bayer pattern) is shown. A cross-section of the sensor is shown along with an example of the spectral sensitivities of the color filters. Left panel adapted from Colin M.L. Burnett/Wikimedia Commons (https://commons.wikimedia.org/wiki/File:-Bayer_pattern_on_sensor.svg) (CC BY-SA 3.0).

the shallower is the DoF. In photography, especially when imaging human subjects for portraits, it is often desirable to have a narrow DoF to focus on the subject's face while blurring out the background. The small aperture used on a mobile camera exhibits little DoF blur. In addition, smartphone cameras have fixed apertures that do not allow for DoF adjustment at capture time. To overcome this limitation, most smartphones now provide a synthetic DoF blur referred to as digital bokeh (see Section 4.3).

2.4. Limited Zoom

As noted above, in response to consumer demands, smartphone design has trended toward ultra-thin form factors. This design trend imposes severe limitations on the thickness (or z-height) of the smartphone's camera module, limiting the effective focal length, which in turn limits the camera module's optical zoom capability. To overcome this z-height limitation, modern smartphone manufacturers typically feature multiple camera modules with different effective focal lengths and fields of view, enabling zoom capabilities ranging from ultrawide to telephoto zoom. The z-height form factor restriction has spurred a so-called thinnovation (a portmanteau of thin and innovation) in optical design, with manufacturers exploring folded optics in an effort to increase the optical path and effective focal length beyond the physical z-height limits of the device.

2.5. Color Subsampling

Finally, a key limitation for both smartphones and most DSLR cameras is that the sensors have only a single color filter associated with each pixel's photodiode,¹ as shown in **Figure 5**. This is analogous to how the human eye's cone cells are categorized by their sensitivity to either short-wavelength, medium-wavelength, or long-wavelength light. Of course, for any camera, the ultimate goal is three color values per pixel. As a result, an interpolation process (called demosaicing) is required to convert the sensor's subsampled color image to one with a three-channel [red, green, and blue (RGB)] value at each pixel. In addition, the RGB color filters used on the camera sensor do not correspond to the perceptual-based CIE XYZ matching functions (Jiang et al. 2013). As

¹Exceptions do exist, including sensors developed by Foveon and others, although these are not in common use.

a result, the ability to produce correct colorimetric measurements is often limited and related to the color filters used.

3. THE CAMERA IMAGING PIPELINE

In this section, we provide an overview of the steps applied by a digital camera to process the image recorded by the sensor and produce the final image that represents a high-quality, visually pleasing photo. These processing steps are applied in sequence and thus form a pipeline where the image pixel values are transformed step by step to produce the final image. Camera systems have a dedicated chip, referred to as an image signal processor (ISP), that performs this processing pipeline in a matter of milliseconds for each image.

We first describe the basic sensor design, followed by a description of a typical single-image camera pipeline. Single-image pipelines are still common on DSLR devices and are sometimes used by mobile phone cameras under good lighting conditions. We then describe multiframe (or burst mode) pipelines that capture and fuse multiple images per photo, more typically used in mobile cameras. Recent advances in multiframe imaging have been instrumental in helping mobile phone cameras overcome many of the limitations described in the previous section.

3.1. Camera Sensor

A camera sensor comprises a 2D grid of photodiodes. A photodiode is a semiconductor device that converts photons (light radiation) into electrical charge. A single photodiode typically corresponds to a single image pixel. To produce a color image, color filters are placed over the photodiodes. These color filters roughly correspond to the long, medium, and short cone cells found in the retina. The typical arrangement of this color filter array (CFA) is often called a Bayer pattern, named after Bryce Bayer, who proposed this design at Kodak in 1975 (Bayer 1975). The CFA appears as a mosaic of color tiles laid on top of the sensor, as shown in **Figure 5**. A key process in the camera pipeline is to demosaic the CFA array by interpolating a RGB value for each pixel based on the surrounding RGB colors. It is important to note that the spectral sensitivities of the red, green, and blue color filters are specific to a particular sensor's make and model. Because of this, a crucial step in the camera imaging pipeline is to convert these sensor-specific RGB values to a device-independent perceptual color space, such as CIE 1931 XYZ. An image captured directly from a sensor that is still in its mosaiced format is called a Bayer image or Bayer frame.

3.2. The Camera Pipeline

Figure 6 shows a diagram of a typical camera imaging pipeline that would be implemented by an ISP. Depending on the ISP design, the routines shown may appear in a slightly different order. Many of the routines described would represent proprietary algorithms specific to a particular ISP manufacturer. Two different camera manufacturers may use the same ISP hardware but can tune and modify the ISP's parameters and algorithms to produce images with a photographic quality unique to their respective devices. This section provides a description of each of the processing steps outlined in **Figure 6**.

3.2.1. Sensor frame acquisition. When the Bayer image from the camera's sensor is captured and passed to the ISP, the ISO gain factor is adjusted at capture time depending on the scene brightness, desired shutter speed, and aperture. The sensor Bayer frame is considered an unprocessed image and is commonly referred to as a raw image. As shown in **Figure 5**, the Bayer frame

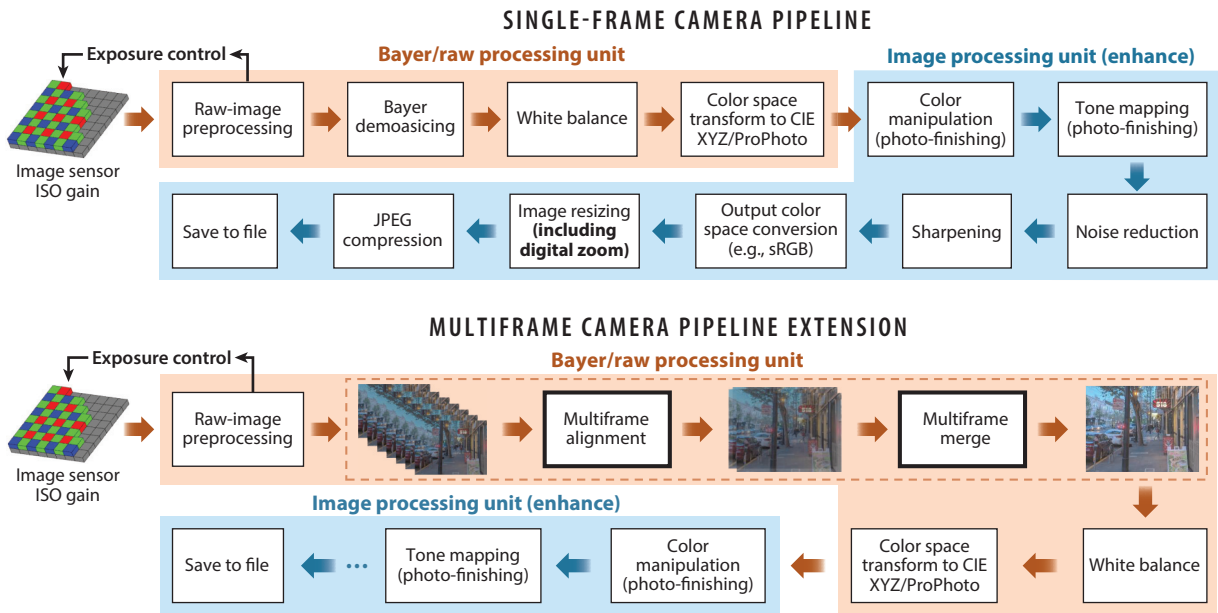


Figure 6

(Top) A standard single-frame camera pipeline. (Bottom) The extension to multiframe (or burst) imaging used by most modern smartphone cameras. Abbreviation: sRGB, standard RGB. Image sensor ISO gain images adapted from Colin M.L. Burnett (CC BY-SA 3.0). Photos adapted from Wronski et al. (2019) (open access).

has a single RGB value per pixel location. These raw RGB values are not in a perceptual color space but are specific to the color filter array’s spectral sensitivities.

3.2.2. Raw-image preprocessing. The raw sensor image is normalized such that its values range from 0 to 1. Many cameras provide a `BlackLevel` parameter that represents the lowest pixel value produced by the sensor. Interestingly, this deviates from 0 due to sensor error. For example, a sensor that is exposed to no light should report a value of 0 for its output, but instead, it outputs a small positive value called the `BlackLevel`. This `BlackLevel` is subtracted off the raw image. The `BlackLevel` is often image specific and related to other camera settings, including ISO and gain. An additional `WhiteLevel` (maximum value) can also be specified. If nothing is provided, the `min` and `max` values of all intensities in the image are used to normalize the image between 0 and 1 after the `BlackLevel` adjustment has been applied.

The preprocessing stage also corrects any defective pixels on the sensor. A defect pixel mask is precalibrated in the factory and marks locations that have malfunctioning photodiodes. Defective pixels can be photodiodes that always report a high value (a hot pixel) or pixels that output no value (a dead pixel). Defective pixel values are interpolated using their neighbors.

Finally, lens-shading (or flat-field) correction is applied to correct the effects of uneven light hitting the sensor. The role of lens-shading correction is shown in **Figure 7**. The figure shows the result of capturing a flat illumination field before lens-shading correction. The amount of light hitting the sensor falls off radially toward the edges. The necessary radial correction is represented as a lens-shading correction mask that is applied by the ISP to correct the effects from the nonuniform fallout. The lens-shading mask is precalibrated by the manufacturer and is adjusted slightly per frame to accommodate different brightness levels, gain factors, and the estimated scene illumination used for white balance (described below).




Figure 7

Light entering the camera does not fall evenly across the sensor. This creates an undesired vignetting effect. Lens-shading correction is used to adjust the recorded values on the sensor to have a uniform response. Figure adapted from Karaimer & Brown (2016).

3.2.3. Bayer demosaicing. A demosaicing algorithm is applied to convert the single-channel raw image to a three-channel full-size RGB image. Demosaicing is performed by interpolating the missing values in the Bayer pattern based on neighboring values in the CFA. **Figure 8** shows an example of the demosaicing process. In this example, a zoomed photodiode with a red color filter is shown. This pixel's green and blue color values need to be estimated. These missing pixel values are estimated by interpolating the missing green pixel using the neighboring green values. A per-pixel weight mask is computed based on the red pixel's similarity to neighboring red pixels. The use of this weight mask in the interpolation helps to avoid blurring around scene edges. **Figure 8** illustrates a simplistic and generic approach; most demosaicing algorithms are proprietary methods that often also perform highlight clipping, sharpening, and some initial denoising (Longere et al. 2002).²

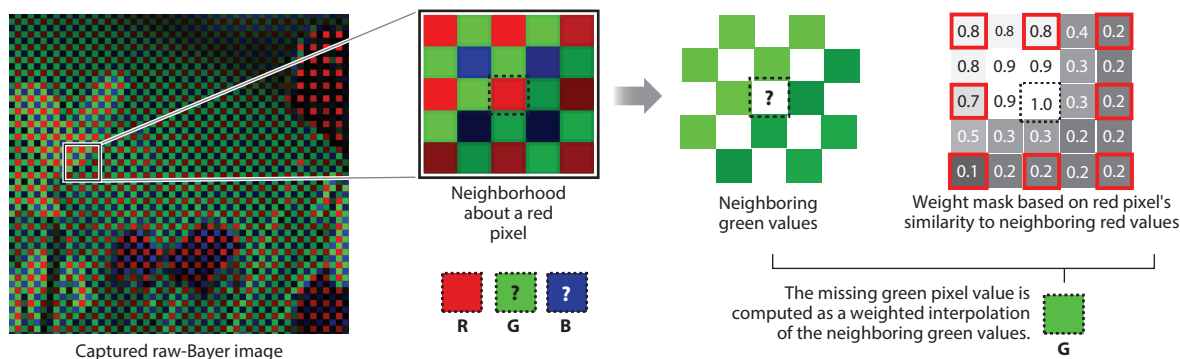


Figure 8

A common approach to image demosaicing. Shown is a red pixel and its neighboring Bayer pixels. The missing green and blue pixel values need to be estimated. These missing values are interpolated from the neighboring pixels. A weight mask based on the red pixel's similarity to its neighbors is computed to guide this interpolation. This weighted interpolation helps to avoid blurring across scene edges. Shown is the interpolation of the missing green pixel value. Raw-Bayer image at left adapted from Cmglee/Wikimedia Commons (https://commons.wikimedia.org/wiki/File:Colorful_spring_garden_Bayer.png) (CC BY-SA 3.0).

²The astute reader will note that this demosaicing step is effectively interpolating two out of three colors at every pixel in the output image. The naive consumer may be shocked to learn that two-thirds of their image is made up.

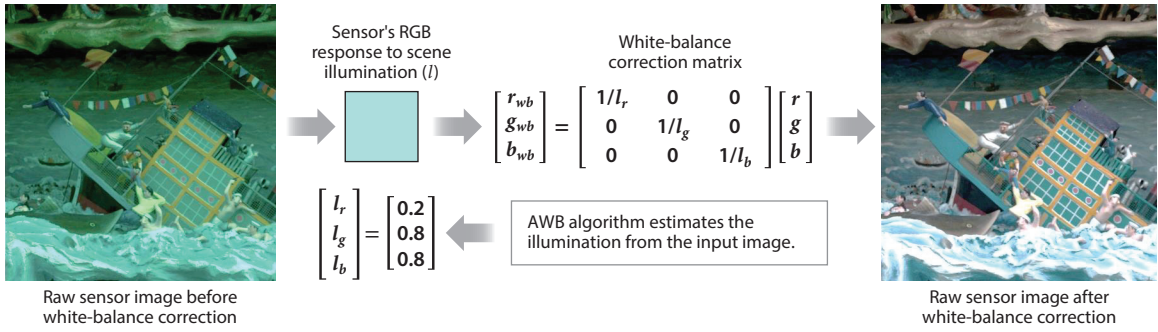


Figure 9

White balance is applied to the image to mimic the human visual system's ability to perform color constancy. An auto-white-balance (AWB) algorithm estimates the sensor's response to the scene illumination. The raw RGB values of the image are then scaled based on the estimated illumination.

3.2.4. White balance. White balance is performed to mimic the human visual system (HVS)'s ability to perform chromatic adaptation to the scene illumination. White balance is often referred to as computational color constancy to denote the connection to the HVS. White balance requires an estimate of the sensor's RGB color filter response to the scene illumination. This response can be precalibrated in the factory by recording the sensor's response to spectra of common illuminations (e.g., sunlight and incandescent and fluorescent lighting). These precalibrated settings then become part of the camera's white-balance preset, which a user can select. A more common alternative is to rely on the camera's auto-white-balance (AWB) algorithm, which estimates the sensor's R, G, B response to the illumination directly from the captured image. Illumination estimation is a well-studied topic in computer vision and image processing, with a wide range of solutions (Barron & Tsai 2017, Buchsbaum 1980, Cheng et al. 2014, 2015, Gehler et al. 2008, Hu et al. 2017, Van De Weijer et al. 2007). **Figure 9** illustrates the white-balance procedure.

Once the sensor's RGB values of the scene illumination have been obtained, either by a preset or by the AWB feature, the image is modified (i.e., white balanced) by dividing all pixels for each color channel by the RGB illumination value. This is similar to the well-known diagonal von Kries color adaption transform (Ramanath & Drew 2014). The von Kries model is based on the response of the eye's short, medium, and long cone cells, while white balance uses the sensor's RGB color filter responses.

3.2.5. Color space transform. After white balance is applied, the image is still in the sensor-specific RGB color space. The color space transform step is performed to convert the image from the sensor's raw-RGB color space to a device-independent perceptual color space derived directly from the CIE 1931 XYZ color space. Most cameras use the wide-gamut ProPhoto RGB color space (Süsstrunk et al. 1999). ProPhoto is able to represent 90% of colors visible to the average human observer.

3.2.6. Color manipulation. Once the image is in a perceptual color space, cameras apply proprietary color manipulation to enhance the visual aesthetics of the image. For DSLR devices, this enhancement can be linked to different picture styles or photo-finishing modes that the user can select, such as vivid, landscape, portrait, and standard. Such color manipulation is often implemented as a 3D lookup table (LUT) that is used to map the input ProPhoto RGB values to new

a Different photo-finishing picture styles



b Photo-finishing

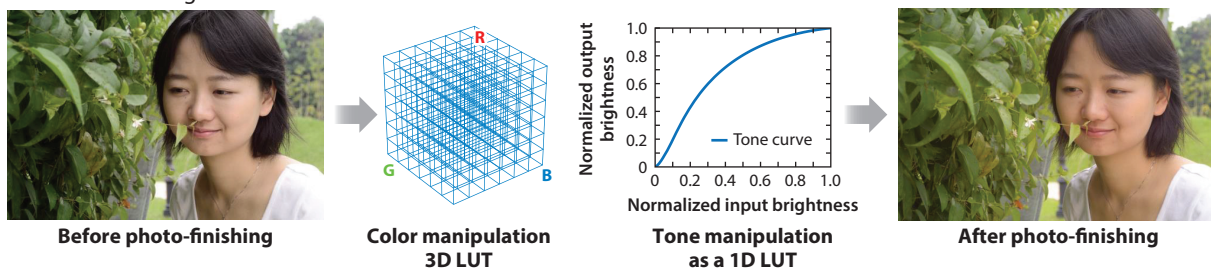


Figure 10

Photo-finishing is used to enhance the aesthetic quality of an image. (a) Cameras often have multiple picture styles. (b) The color manipulation is often performed as a combination of a 3D lookup table (LUT) to modify the RGB colors and a 1D lookup table to adjust the image's tonal values. 3D LUT image adapted from Karaimer & Brown (2016).

RGB values based on a desired manipulation. **Figure 10** shows an example. A 1D LUT tone map (discussed in the next section) is also part of this photo-finishing manipulation.

Additional color manipulation may be performed on a smaller set of select colors used to enhance skin tones. Establishing the 3D LUT can be a time-consuming process and is often performed by a group of golden-eye experts who tune the ISP algorithms and tables to produce a particular photographic aesthetic often associated with a particular camera. Note that camera manufacturers may even sell the same camera with different color manipulation parameters based on the preferences of users in different geographical locations. For example, cameras sold in Asia and South America often have a slightly more vivid look than those sold in European and North American markets.

3.2.7. Tone mapping. A tone map is a 1D LUT that is applied per color channel to adjust the tonal values of the image. **Figure 10** shows an example. Tone mapping serves two purposes. First, combined with color manipulation, it adjusts the image's aesthetic appeal, often by increasing the contrast. Second, the final output image is usually only 8 to 10 bits per channel (i.e., 256 or 1,024 tonal values), while the raw-RGB sensor represents a pixel's digital value using 10–14 bits (i.e., 1,024 up to 16,384 tonal values). As a result, it is necessary to compress the tonal values from the wider tonal range to a tighter range via tone mapping. This adjustment is reminiscent of the human eye's adaptation to scene brightness (Land 1974). **Figure 11** shows a typical 1D LUT used for tone mapping.

3.2.8. Noise reduction. Applying noise reduction algorithms is a key step to improving the visual quality of the image. A delicate balance must be struck in removing image noise while avoiding the suppression of fine-detail image content. Too-aggressive denoising may give the image a




Figure 11

An example of a global tone map applied to a high-dynamic-range (HDR) image and local tone maps that vary spatially based on the image content. Photos adapted from Afifi et al. (2019).

blurred appearance. Too little image denoising may result in visual noise being dominant and distracting in the final image. Given the importance of denoising, there is a large body of literature on this problem, which we discuss in detail in Section 3.4.1. Denoising algorithms often consider multiple factors, including the captured image’s ISO gain level and exposure settings. While we show noise reduction applied after the color space transform and photo-finishing, some ISPs apply noise reduction before photo-finishing, or both before and after. Indeed, ISPs often provide denoising algorithms that can be tuned by camera makers to taste.

3.2.9. Output color space conversion. At this stage in the camera pipeline, the image’s RGB values are in a wide-gamut ProPhoto color space. However, modern display devices can only produce a rather limited range of colors. As a result, the image is converted to a display-referred (or output-referred) color space intended for consumer display devices with a narrow color gamut. The most common color space is the standard RGB (sRGB), which is also the standard for HDTV. Other color spaces, such as AdobeRGB and Display-P3, are sometimes used. The output color space conversion includes a tone-mapping operator as part of its color space definition. This final tone-mapping operator is referred to as a gamma encoding. The name comes from the Greek letter used in the formula to model the nonlinear tone curve. The purpose of the gamma encoding is to code the digital values of the image into a perceptually uniform domain (Poynton 2012). The gamma values used for sRGB and Display-P3 closely follow Stevens’s power law coefficients for perceived brightness (Stevens 1961).

3.2.10. Image resizing. The image can be resized based on the user preferences or target output device (e.g., if the camera is used in a preview mode with a viewfinder). Image resizing is not limited to image downsizing and can be employed to upsample a cropped region in the captured image to a larger size to provide a digital zoom. More details of this operation appear in Section 4.2.2.

3.2.11. JPEG compression and metadata. The image is finally compressed, typically with the JPEG compression standard, and saved. Additional information, such as capture time, GPS location, and exposure setting, can be saved with the image as metadata.

3.3. Modern Multiframe (Burst) Pipeline

The advancement of smartphone displays and networking capabilities in 2010, together with the emergence of image-sharing services like Instagram and Pinterest, resulted in users taking many more photographs with their smartphones and sharing them more broadly. Seeing their photographs mixed alongside professionally produced content for the first time, users began




Figure 12

Burst photography used in most mobile imaging pipelines consists of four major steps. In the first, a burst of frames is captured based on an exposure schedule defining the total number of frames to capture as well as the exposure time for each individual frame. This defines the total exposure time for the burst. In the second, frames in the burst are spatially aligned. In the third, the aligned frames are merged into a single output frame. The fourth step, enhancement, encapsulates all postprocessing steps after the merge step, including color manipulation, tone mapping, and noise reduction. Figure adapted from Wronski et al. (2019) (open access).

demanding higher quality from their smartphone cameras. This increased demand for higher quality spurred innovation in the industry. Smartphone manufacturers began utilizing the additional computing power of the smartphone by incorporating computational photography techniques in an effort to overcome the limitations enforced by the smartphone's physical form and to bridge the quality gap relative to dedicated camera devices, like DSLR cameras.

Over the course of the past decade, the modern smartphone camera pipeline has evolved around the concept of capturing and merging multiple frames (known as burst processing) to generate images of greater quality than is possible through the capture of a single image alone. This approach has a comparatively long history in the computational photography research domain, where multiframe processes have been proposed for denoising (Godard et al. 2018, Mildenhall et al. 2018, Telleen et al. 2007), joint demosaicing and denoising (Gharbi et al. 2016, Tan et al. 2017), joint demosaicing and super-resolution (Farsiu et al. 2006), and high-dynamic-range imaging (Debevec & Malik 1997, Mertens et al. 2007). Recent years have seen the incorporation of multiframe techniques like these into the default photography mode of smartphone cameras aimed at synthesizing higher-resolution sensors with larger pixels and higher bit depth.

To implement burst processing, smartphones can incorporate a strategy of continuous capture, where, on launching the camera, frames are continuously captured and stored into a ring buffer.³ In this mode, known as zero shutter lag (ZSL), when the shutter is pressed, the buffer of frames is passed to the camera processing pipeline for merging. The merge process selects a single frame from the burst close to the shutter press as a base frame and aligns and merges information from surrounding frames to improve image quality. Two critical factors in this process are the correct exposure of captured frames to ensure adequate light capture with minimal motion blur and accurate motion alignment between frames. In this way, the processing pipeline aims to loosely emulate our own ability to annul image motion by adaptation of our spatiotemporal receptive fields (Burr et al. 1986). The general structure shown in **Figure 12** illustrates the modern burst processing pipeline used in many mobile cameras, consisting of exposure control, frame alignment, merging, and postmerge image enhancement components. We explore these components below.

³A ring buffer stores a sequence of frames in order and has the property that, when it is full and a subsequent new frame is captured, the new frame overwrites the oldest frame in the buffer, keeping a constant number of frames in memory.

3.3.1. Exposure control. Collecting sufficient light through accurate exposure is critical for controlling image noise, blur, and dynamic range for high-quality photography (Hasinoff et al. 2010). The signal-to-noise ratio varies proportionally to the exposure time, so for shadow regions or dimly lit scenes, it is important to set the exposure high enough to suppress noise and capture sufficient image detail. However, setting the exposure too high can cause the captured light to exceed the image sensor’s pixel well capacity, resulting in saturated (or blown-out) image details.

Exposure can be adjusted through the combination of three camera settings: the aperture, the ISO (sensor gain or effective sensor sensitivity), and the shutter speed (exposure time). However, given that most smartphones have a fixed aperture, exposure control is typically limited to adjustment of the sensor gain and exposure time. Some smartphone cameras do provide manual control over the exposure, but most target the nonexpert user and instead control the exposure automatically.

Approaches to estimating the correct exposure time for a single captured image often use measurements such as the distance between the maximum and minimum luminance, as well as the average luminance in the scene (Sampat et al. 1999). In a smartphone’s burst processing pipeline, however, exposure control is more complex, since the captured image is generated by merging multiple frames. For burst processing, the camera must define a schedule of exposure times for a burst of captured frames to achieve an overall total exposure time for the merged output frame (see **Figure 12**). For example, a total exposure time of 300 ms could be achieved through a schedule of five frames, each with an exposure time of 60 ms. Similarly, in a burst processing pipeline implementing high-dynamic-range (HDR) through bracketing, the exposure schedule might define a schedule of short, medium, and long exposures. Exposure control for burst processing therefore needs to take into consideration not only the available light in the scene but also the merge processing and how it impacts the overall exposure.

Another factor that greatly impacts exposure control is camera shake (e.g., due to being hand-held), which can introduce motion blur. To enable longer exposure times, modern smartphone cameras incorporate an optical image stabilizer (OIS), which actively counteracts camera shake. However, this often does not completely remove the motion and does not help in the case of (local) subject motion, which is also a source of motion blur. Adapting the exposure schedule in accordance with motion observed in the scene is a common approach used to reduce the impact of motion blur. In Section 4.1, we further examine exposure control and more severe limitations in the case of low-light photography.

3.3.2. Alignment. Generating high-quality, artifact-free images through burst processing relies on robust and accurate spatial alignment of frames in the captured burst. This alignment process must account for not only global camera motion (residual motion not compensated for by the OIS) but also local motion in the scene. There is a long history of frame alignment techniques in the research literature, from early variational methods that solve the global alignment problem using assumptions of brightness constancy and spatial smoothness (Horn & Schunck 1981) to multiscale approaches solving for both global and local motion (Bruhn et al. 2005).

Given the omnipresent convenience of the smartphone, photography has been made possible in the most extreme of conditions. As a result, accurate alignment can be challenging, particularly in underexposed or low-light scenes, where noise can dominate the signal. Similarly, overexposed scenes can introduce clipping or motion blur, making alignment difficult or impossible due to the loss of image detail. Even with optimal exposure, complex nonrigid motion, lighting changes, and occlusions can make alignment challenging.

Although state-of-the-art multiscale deep learning methods can achieve accurate frame alignment in challenging conditions (Sun et al. 2018), they are currently beyond the computational

capabilities of many smartphones. As a result, burst processing on a smartphone is greatly limited by the accuracy of the alignment process. The exposure schedule must be defined so as to facilitate accurate alignment, and the merge method must in turn be designed to be robust to misalignments to avoid jarring artifacts in the merged output (also known as fusion artifacts). Common merge artifacts due to misalignments include ghosting and zipper artifacts, often observed along the edges of moving objects in a captured scene.

3.3.3. Merge. Once accurately aligned, the smartphone's burst processing pipeline must reduce the multiple captured frames to a single output frame. In static scenes and in the absence of camera shake, a simple averaging of frames of the same exposure will reduce noise proportionally to the square root of the total number of merged frames. However, few scenarios arise in real-world photography where such a simple merging strategy could be effectively applied. In addition, such a simple merge strategy underutilizes the burst processing approach, which, as mentioned above, can also facilitate increasing dynamic range and resolution. In this section, we describe a merge method aimed at reducing noise and increasing dynamic range called HDR+ (Hasinoff et al. 2016); in Section 4.2, we describe a generalization of the method aimed at increasing resolution as well.

HDR+ was one of the earliest burst processing approaches that saw mass commercial distribution, featured in the native camera app of Google's Nexus and Pixel smartphones. Aimed at reducing noise and increasing dynamic range, the HDR+ method employs a robust multiframe merge process operating on 2–8 frames, achieving interactive end-to-end processing rates. To reduce the impact of motion blur and to avoid pixel saturation, HDR+ defines an exposure schedule to deliberately underexpose the captured frames in a ZSL buffer. Bypassing the smartphone's standard (single-frame) ISP, the merge pipeline operates on the raw-Bayer frames directly from the camera's sensor, enabling the merge process to benefit from higher bit-depth accuracy and simplifying the modeling of noise in the pipeline.

Given a reference frame close (in time) to the shutter press, the HDR+ pipeline successively aligns and merges alternative frames in the burst in a pair-wise manner. The merging of frame content operates on tiles and is implemented in the frequency domain. For each pair of reference and alternate tiles, a new tile is linearly interpolated between them (per frequency) and averaged with the reference tile to generate the merged tile output. Given that the merge strategy is applied per frequency, the merging achieved per tile can be partial. The interpolation weight is defined as a function of the measured difference between the aligned tile pairs and the expected (i.e., modeled) noise. For very large measured differences (e.g., possibly due to misalignment), the interpolated output tends toward the reference tile, whereas for differences much less than the expected noise, the interpolated output tends toward the alternate tile. By adapting in this way, the merging process provides some robustness to misalignment and degrades gracefully to outputting the reference frame only in cases where misalignment occurs across the entire burst.

The final output of the merge process is a Bayer frame with higher bit depth and overall SNR, which is then passed to a postmerge enhancement stage, including demosaicing, color correction, and photo-finishing. Of particular importance among these postmerge processes is spatial denoising. As a consequence of the tile-wise and partial merging of frames, the resulting merged frame can have spatially varying noise strength, which must be adequately handled by the postmerge denoising process.

3.4. Photo-Finishing: Denoising, Tone Mapping, and Sharpening

While significantly improved quality results from merging multiple frames, it is nevertheless the case that, just as in the single-frame pipeline, several additional photo-finishing operations are still

required to give the final picture the desired aesthetic qualities of a pleasing photograph. These steps (denoising, tone mapping, and sharpening) are some of the core technologies of traditional image processing and have been studied extensively over the several decades preceding the advent of what we now call computational photography. However, with the heightened demands of new sensors on smaller devices, the development of these (single-image) enhancement techniques too has been accelerated in recent years. Indeed, the multiframe approach may require aspects of photo-finishing to be tailored for the multiframe merged output. As discussed in Section 2, denoising is an operation of great importance in the pipeline (whether single or multiframe), and thus we begin by describing this operation in some detail.

3.4.1. Denoising. As should be clear from the discussion above, filtering an image is a fundamental operation throughout the computational photography pipeline. The most widely used canonical filtering operation is one that smooths an image or, more specifically, removes or attenuates the effect of noise.

The basic design and analysis of image denoising operations have informed a very large part of the image processing literature (Lebrun et al. 2012, Milanfar 2013, Takeda et al. 2006, 2007), and the resulting techniques have often quickly spread or been generalized to address a wider range of restoration and reconstruction problems in imaging.

Over the past five decades, many approaches have been tried, starting with the simplest averaging filters and moving to filters that adapt somewhat better (but still rather empirically) to the content of the given image. With shrinking device sizes, and the rise in the number of pixels per unit area of the sensor, modern mobile cameras have become increasingly prone to noise. The manufacturers of these devices, therefore, depend heavily on image denoising algorithms to reduce the spurious effects of noise.

Only in the past decade or so (and concomitant with the broad proliferation of mobile devices) has a great leap forward in denoising performance been realized (Chatterjee & Milanfar 2010). This recent progress was ignited by patch-based methods (Buades et al. 2005, Efros & Leung 1999). This generation of algorithms exploits both local and nonlocal redundancies or self-similarities in the image. The now-commonplace procedure is to measure and make use of affinities (or similarities) between a given pixel or patch of interest and other pixels or patches in the given image. These similarities are then used in a filtering (e.g., data-dependent weighted-averaging) context to give higher weights to contributions from more similar data values and to properly discount data points that are less similar.

The bilateral filter (Tomasi & Manduchi 1998) was developed with this idea in mind, as were its spiritually close predecessors like the Susan filter (Smith & Brady 1997). More recent extensions of these ideas include the work of Dabov et al. (2007), Buades et al. (2005), Takeda et al. (2006), and Zoran & Weiss (2011), and other generalizations described by Milanfar (2013).

The general construction of many denoising filters begins by specifying a (symmetric positive semidefinite) kernel $k_{ij}(\mathbf{y}) = K(y_i, y_j) \geq 0$, from which the coefficients of the adaptive filter are constructed. In this case, \mathbf{y} denotes the noisy image, and y_i and y_j denote pixel values at locations i and j , respectively.⁴

⁴In practice, it is commonplace to compute the kernel not on the original noisy \mathbf{y} , but on a prefiltered version of it, processed with some basic smoothing, with the intent of weakening the dependence of the filter coefficients on noise.

Specifically, we have

$$w_{ij} = \frac{k_{ij}}{\sum_{i=1}^n k_{ij}}.$$

Each pixel of the denoised image $\hat{\mathbf{y}}$ is then given by

$$\hat{y}_j = \sum_{i=1}^n w_{ij} y_i,$$

where the coefficients $[a_{1j}, \dots, a_{nj}]$ describe the relative contribution of the input (noisy) pixels to the output pixels, with the constraint that they sum to one:

$$\sum_{i=1}^n a_{ij} = 1.$$

Below, we concretely highlight a few such kernels that lead to popular denoising or smoothing filters. These filters are commonly used in the computational photography, imaging, computer vision, and graphics literature for many purposes.

3.4.2. Classical Gaussian filters. Measuring (only) the spatial distance between two pixels located at (2D) spatial positions \mathbf{x}_i and \mathbf{x}_j , the classical Gaussian kernel is

$$k_{ij} = \exp\left(\frac{-\|\mathbf{x}_i - \mathbf{x}_j\|^2}{b^2}\right).$$

Such kernels lead to the classical and well-worn Gaussian filters that apply the same weights regardless of the underlying pixels.

3.4.3. The bilateral filter. The bilateral filter takes into account both the spatial and the data-wise distances between two samples, in separable fashion, per Tomasi & Manduchi (1998) and Elad (2002):

$$k_{ij} = \exp\left(\frac{-\|\mathbf{x}_i - \mathbf{x}_j\|^2}{b_x^2}\right) \exp\left(\frac{-(y_i - y_j)^2}{b_y^2}\right) = \exp\left\{\frac{-\|\mathbf{x}_i - \mathbf{x}_j\|^2}{b_x^2} + \frac{-(y_i - y_j)^2}{b_y^2}\right\}.$$

As can be observed in the exponent on the right-hand side, the similarity metric is a weighted Euclidean distance between the vectors (\mathbf{x}_i, y_i) and (\mathbf{x}_j, y_j) . This approach has several advantages. Namely, while the kernel is easy to construct, and computationally simple to calculate, it yields useful local adaptivity to the pixels.

3.4.4. Nonlocal means. The nonlocal means algorithm, originally proposed by Buades et al. (2005), is a generalization of the bilateral filter in which the data-dependent distance term (see the equation in the previous section) is measured block-wise instead of point-wise:

$$k_{ij} = \exp\left(\frac{-\|\mathbf{x}_i - \mathbf{x}_j\|^2}{b_x^2}\right) \exp\left(\frac{-\|\mathbf{y}_i - \mathbf{y}_j\|^2}{b_y^2}\right),$$

where \mathbf{y}_i and \mathbf{y}_j refer to patches (subsets of pixels) in \mathbf{y} .

3.4.5. Locally adaptive regression kernel. The key idea behind this kernel is to robustly measure the local structure of data by making use of an estimate of the local geodesic distance between nearby samples (Takeda et al. 2007):

$$k_{ij} = \exp \{ -(\mathbf{x}_i - \mathbf{x}_j)^T \mathbf{Q}_{ij} (\mathbf{x}_i - \mathbf{x}_j) \},$$

where $\mathbf{Q}_{ij} = \mathbf{Q}(\mathbf{y}_i, \mathbf{y}_j)$ is the covariance matrix of the gradient of sample values estimated from the given pixels, yielding an approximation of local geodesic distance in the exponent. The dependence of \mathbf{Q}_{ij} on the given data means that the denoiser is highly nonlinear and shift varying. This kernel is closely related to, but somewhat more general than, the Beltrami kernel of Spira et al. (2007) and the coherence-enhancing diffusion approach of Weickert (1999).

More recently, methods based on deep convolutional neural networks (CNNs) have become dominant in terms of the quality of the overall results with respect to well-established quantitative metrics (Burger et al. 2012, Meinhardt et al. 2017). Supervised deep-learning-based methods are currently the state of the art (for example, Burger et al. 2012, Chen et al. 2015, D. Liu et al. 2018, P. Liu et al. 2018, Mao et al. 2016, Remez et al. 2018, Tai et al. 2017, Wang et al. 2015, Zhang et al. 2017, K. Zhang et al. 2018, Y. Zhang et al. 2018). However, these CNN-based approaches have yet to become practical (especially in mobile devices) due not only to heavy computation and memory demands, but also to their tendency to sometimes also produce artifacts that are unrealistic with respect to more qualitative perceptual measures.

Finally, it is worth noting that, as denoising methods have evolved from traditional signal- or image-processing approaches to deep neural networks (DNNs), there has been an increasing need for training sets comprised of images that accurately represent noise found on small sensors used in camera phones. The recent DND (Plötz & Roth 2017) and SSID data sets (Abdelhamed et al. 2018) provide images captured directly from such devices for use in DNN training. Both data sets have shown that training using real images versus those synthesized from existing noise models provides improved performance. This indicates the need for better noise models in the literature that are able to capture the true characteristics of small camera sensors.

3.4.6. Tone mapping. As discussed in Section 3.2, tone mapping is applied as part of the photo-finishing routines in both single-frame and multiframe camera pipelines. Tone mapping manipulates the intensity levels, or tones, in the image. Assume that the variable r represents the intensity levels for one of the color channels (RGB) that will be enhanced. Assume also that the image intensity values have been normalized such that they lie on the interval $[0, 1]$. A tone map can be expressed as follows:

$$s = T(r),$$

where function T produces a new intensity level s for every input level r (i.e., it maps input tones to output tones). Tone mapping is applied either globally or locally. A global tone map, often referred to as a tone curve, is applied to all pixels' intensity values in the image irrespective of the pixel's spatial location. A global tone map is constrained to satisfy the following conditions: (a) $T(r)$ is single-valued and monotonically increasing, and (b) $0 \leq T(r) \leq 1$.

Because images have discrete intensity values, a global $T(r)$ can be implemented as a 1D LUT. A global tone map can be applied to each color channel, or a separate tone map can be designed per color channel. In addition, tone maps can be customized depending on the mode of imaging. For example, in burst mode for low-light imaging, a tone map can be adjusted to impart a night scene's look and feel. This can be done using a tone map that maintains strong contrast with dark shadows and strong highlights (Levoy & Pritch 2018).

Local tone mapping adjusts intensity values in a spatially varying manner, inspired by the HVS's sensitivity to local contrast. Local tone mapping methods, often referred to as tone operators, examine a spatial neighborhood about a pixel to adjust the intensity value (Ahn et al. 2013, Banterle et al. 2017, Cerda et al. 2018, Ma et al. 2015, Paris et al. 2011, Reinhard et al. 2002). As a result, the single-valued and monotonicity constraints are not always enforced, as is done in global tone mapping. For example, in the case of burst imaging for HDR, intensity values from the multiple images can be combined in a manner that darkens or lightens regions to enhance the image's visual quality. Most methods decompose the input image into a base layer and one or more detail layers. The detail layers are adjusted based on local contrast, while a global tone map modifies the base layer. **Figure 11** shows an HDR image that has been processed using both global and local tone mapping.

3.4.7. Sharpening. Image sharpness is one of the most important attributes that defines the visual quality of a photograph. Every image processing pipeline has a dedicated component to mitigate the blurriness in the captured image. Although there are several methods that try to quantify the sharpness of a digital image, there is no clear definition that perfectly correlates with the quality perceived by our visual system. This makes it particularly difficult to develop algorithms and properly adjust their parameters in such a way that they produce appealing visual results in the universe of use cases and introduce minimal artifacts.

Image blur can be observed when the camera's focus is not correctly adjusted, when the objects in the scene appear at different depths, or when the relative motion between the camera and the scene is faster than the shutter speed (motion blur, camera shake). Even when a photograph is perfectly shot, there are unavoidable physical limitations that introduce blur. Light diffraction due to the finite lens aperture, integration of the light in the sensor, and other possible lens aberrations introduce blur, leading to a loss of details. Additionally, other components of the image processing pipeline itself, particularly demosaicing and denoising, introduce blur.

A powerful yet simple model of blur is to assume that the blurry image is formed by the local average of nearby pixels of a latent unknown sharp image that we would like to estimate. This local average acts as a low-pass filter attenuating the high-frequency image content, introducing blur. This can be formally stated as a convolution operation—that is,

$$\mathbf{v}[i, j] = \sum_{k, l} b[k, l] \mathbf{u}[i - k, j - l],$$

where \mathbf{v} is the blurry image that we want to enhance, \mathbf{u} is the ideal sharp image that we do not have access to, and b models the typically unknown blur filter.

There are two conceptually different approaches to removing image blur and increasing apparent sharpness. Sharpening algorithms seek to directly boost high- and mid-frequency content (e.g., image details, image edges) without explicitly modeling the blurring process. These methods are sometimes also known as edge enhancement algorithms since they mainly increase edge contrast. In contrast, deconvolution methods try to explicitly model the blurring process by estimating a blur kernel b and then trying to invert it. In practice, there are an infinite number of possible combinations of \mathbf{u} and b that can lead to the same image \mathbf{v} , which implies that recovering \mathbf{u} from \mathbf{v} is an ill-posed problem. One of the infinite possible solutions is indeed the no-blur explanation: $\mathbf{u} = \mathbf{v}$, and b is the trivial kernel that maintains the image unaltered. This implies that the degradation model is not sufficient to disentangle the blur b and the image \mathbf{u} from the input image \mathbf{v} , and more information about b and/or \mathbf{u} (prior) is needed.

Most blind deconvolution methods proceed in two steps: A blur kernel is first estimated, and then, using the estimated kernel, a nonblind deconvolution step is applied. These methods



Figure 13

Example of deblurring a mildly blurry image using Polyblur. Images adapted from Delbracio et al. (2020). The estimated blur kernel is shown on the top right in the right panel.

generally combine natural image priors (i.e., what characteristics does a natural sharp image have) and assumptions about the blur kernel (e.g., maximum size) to cast the blind deconvolution problem as one of variational optimization (El-Henawy et al. 2014, Fergus et al. 2006, Levin et al. 2009). In the specific case of deblurring slightly blurry images, we can proceed in a more direct way by filtering the image with an estimate of the blur and thus avoid using costly optimization procedures (Delbracio et al. 2020, Hosseini & Plataniotis 2019). **Figure 13** shows an example using Polyblur (Delbracio et al. 2020), which efficiently removes blur by estimating the blur and combining multiple applications of the estimated blur to approximate its inverse.

One of the most popular and simplest sharpening algorithms is unsharp masking. First, a copy of the given image is further blurred to remove high-frequency image details. This new image is subtracted from the original to create a residual image that contains only image details. Finally, a fraction of this residual image is added back to the original, which results in boosting of the high-frequency details. This procedure has its roots in analog photography, where a blurry positive image is combined with the original negative to create a new, more contrasted photograph. A typical digital implementation of unsharp masking is by using a Gaussian blur,

$$\mathbf{u}_{\text{unsharp}} = \mathbf{u} + \kappa(\mathbf{u} - G_{\sigma}\mathbf{u}),$$

where G_{σ} is the Gaussian blur operator having strength σ . The parameter κ and the amount of blur σ should be empirically set.

At a very broad level, the HVS behaves similarly to unsharp masking and the Laplace operator (Ratliff 1965). The center-surround receptive fields present in the eye have both excitatory (center) and inhibitory (surrounding) regions. This leads our visual system to enhance changes in contrast (e.g., edge detection) by exciting and inhibiting regions in a way similar to the action of the Laplace operator. In fact, one manifestation of this phenomenon is the well-known Mach bands illusion, where the contrast between edges is exaggerated by the HVS (**Figure 14**).

There are different variants of this high-frequency boosting principle. Kovácsznay & Joseph (1955) introduced the idea that a mildly blurred image could be deblurred by subtracting a small amount of its Laplacian:

$$\hat{\mathbf{u}} = \mathbf{u} - \kappa \Delta \mathbf{u}.$$

In fact, this method is closely related to unsharp masking, where the residual mask $\mathbf{u} - G_{\sigma}\mathbf{u}$ is replaced with the negative of the image Laplacian $-\Delta \mathbf{u}$.

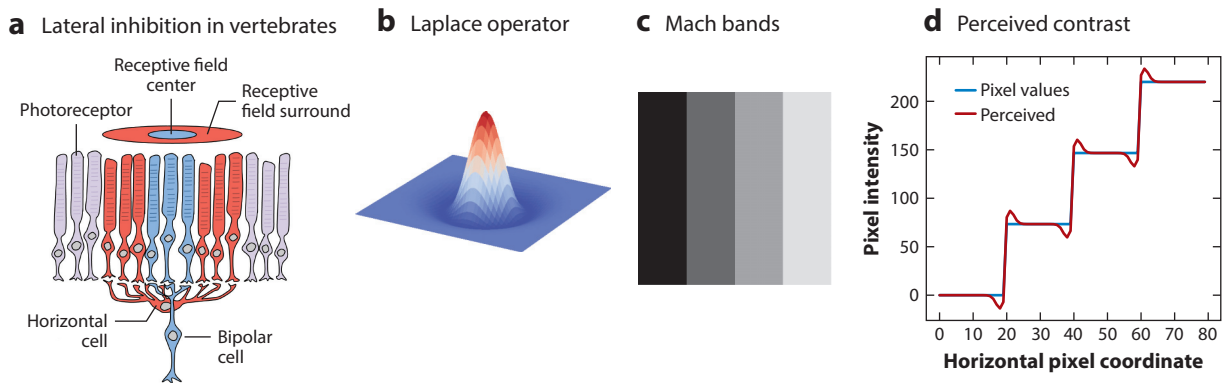


Figure 14

Mach bands. (a) The center-surround receptive fields present in the human visual system have both excitatory (center) and inhibitory (surrounding) regions. Panel adapted from Kramer & Davenport (2015) (CC BY 4.0). (b) This leads our visual system to enhance changes in contrast by exciting and inhibiting regions in a way similar to the action of the Laplace operator. (c) The Mach bands illusion is one manifestation of this phenomenon, where (d) the contrast between edges is exaggerated by the our human visual system.

A perhaps not very well-known fact is that the Nobel Prize winner Dennis Gabor studied this process and determined how to set the best amount to subtract (Lindenbaum et al. 1994). In fact, Lindenbaum et al. (1994) showed that Laplacian sharpening methods can be interpreted as approximating inverse diffusion processes—for example, by diffusion according to the heat equation, but in reverse time. This connection has led to numerous other sharpening methods in the form of regularized partial differential equations (Buades et al. 2006, Osher & Rudin 1990, You et al. 1996).

4. COMPOUND FEATURES

A common adage has emerged that aptly describes the age of mobile cameras: The best camera is the one that is with you. This sentiment expresses that there is no longer any need to carry a second camera if you have a mobile phone in your pocket that has a camera with nearly the same functionality. Of course, the key caveat is nearly the same. Some of what we take for granted in a large-form-factor camera is possible only because of the form factor. To approximate these functions, we must combine various pieces of technology to emulate the end result.

In this section, we briefly describe how we combine the basic techniques described above to enable advanced features that not only approximate some functionalities of larger cameras, but also sometimes even exceed them. For instance, hybrid optical–digital zoom requires state-of-the-art multiframe merge and single-frame upscaling technologies. Another example is synthetic bokeh (e.g., synthesizing shallow DoF), which requires both segmentation of the image for depth and application of different processing to the foreground versus the background.

4.1. Low-Light Imaging

When photographing a low-light or night scene, the goal is often not to capture exactly what we see, but instead to create a visually pleasing image that also conveys the darkness of the scene. Therefore, unlike human vision, which becomes scotopic in dim light with limited color perception (Kelber et al. 2017), smartphone cameras aim to produce photos that are colorful and noise free. Until relatively recently, high-quality photography in very low-light conditions

was achievable only on standalone cameras like DSLR cameras, with large enough pixels and adjustable aperture to enable sufficient light capture. As described in Section 3.3, although the exposure time can be increased synthetically by merging multiple frames, there are other factors that inherently limit the maximum achievable exposure time. In a ZSL shutter mode (see Section 3.3), the frames acquired by the camera pipeline are also used to drive the viewfinder display. To avoid noticeable judder, the viewfinder must achieve a frame rate of at least 15 frames per second directly, limiting the maximum exposure to 66 ms, which is often insufficient for very low-light scenes.

To overcome this, smartphones have adopted a new frame capturing strategy known as positive shutter lag (PSL) in which frames are captured after the shutter press (Levoy & Pritch 2018). By capturing frames after the shutter press, these so-called night modes can achieve exposure times well beyond the previous 66 ms limit. However, to operate robustly in the wild, the camera still needs to automatically adapt to the amount of local and global motion in the scene to avoid the introduction of motion blur. The Night Sight feature of the Pixel smartphone (Liba et al. 2019) solves the motion blur problem through real-time temporal motion metering that runs prior to the shutter press, predicting future scene motion and automatically setting the exposure time and gain (ISO) accordingly, enabling frame exposure times over 300 ms in cases where there is little motion.

In addition to the problem of capturing sufficient light, very low-light conditions make tasks such as AWB challenging. Liba et al. (2019) addressed this through a learning-based approach to AWB that is trained specifically on night scenes. Liba et al. (2019) also introduced a creative tone-mapping solution that draws inspiration from artists' portrayals of night scenes in paintings by keeping shadows close to black and boosting color saturation (Levoy & Pritch 2018).

Extending the low-light capabilities of photography even further, some smartphones now offer astrophotography modes. This has been made possible through more sophisticated motion detection modes that utilize on-device sensors to detect tripod (or nonhandheld) mounting, enabling synthetic exposure times of over four minutes (Kainz & Murthy 2019).

4.2. Super-Resolution and Hybrid Optical and Digital Zoom

Due to the physical limitations on the camera's form factor, one of the principal limitations of a mobile camera is its ability to zoom and focus across a broad range of magnification factors. The thinness of the device prevents the placement of a lens with a broadly variable focal length in front of the sensor. As such, the ability to resolve objects in the mid- and far ranges is inherently limited.

To address these limitations, two broad classes of approaches have been developed. First, complex optical hardware designs such as the periscope lens (Lovejoy 2020) have enabled larger focal length to be implemented inside typically thin mobile devices. These innovative designs have enabled true optical magnifications as high as 5–10 \times to be implemented. However, the focal length of such telephoto lenses is fixed, and the number of such optical elements is still limited to at most one or two due to the scarcity of space inside the phone and mechanical limitations. As a result, almost regardless of the optical power of the elements available, the overall zoom pipeline inside all mobile devices has necessarily evolved as a hybrid of optical and digital magnification techniques.

4.2.1. Multiframe super-resolution. An example of a modern mobile zoom pipeline is the Super Res Zoom pipeline described by Wronski et al. (2019) and used in Google's Pixel devices, illustrated in **Figure 15**. This hybrid optical–digital zoom pipeline first implements a burst processing pipeline that achieves multiframe super-resolution by aligning, merging, and enhancing a sequence of raw frames with subpixel accuracy. This process circumvents the need for the




Figure 15

The full Super Res Zoom pipeline enhances image resolution in two distinct ways: the merge step and the single-frame upscaling step. Figure adapted from Wronski et al. (2019) (open access).

typical (first) demosaicing step described in Section 3.2. As a result, the high-frequency (and often aliased) information in the raw frames is used directly in the formation of a full-resolution image at or near the native resolution of the camera.

This approach implements demosaicing and super-resolution simultaneously, formulating the problem as the reconstruction and interpolation of a continuous signal from a set of possibly sparse samples. Namely, the RGB pixels measured individually on each frame are reconstructed simultaneously onto a common grid. This technique enables the production of highly detailed images, containing information that would have already been lost in part due to the earlier (and much more naive) interpolation in the demosaicing step. An additional advantage is that it allows us to directly create an image with a desired target magnification or zoom factor.

The approach of Wronski et al. (2019) is visualized in **Figure 16**. Similar to the standard merge process described above, first, a burst of raw (CFA Bayer) images is captured. Each captured frame

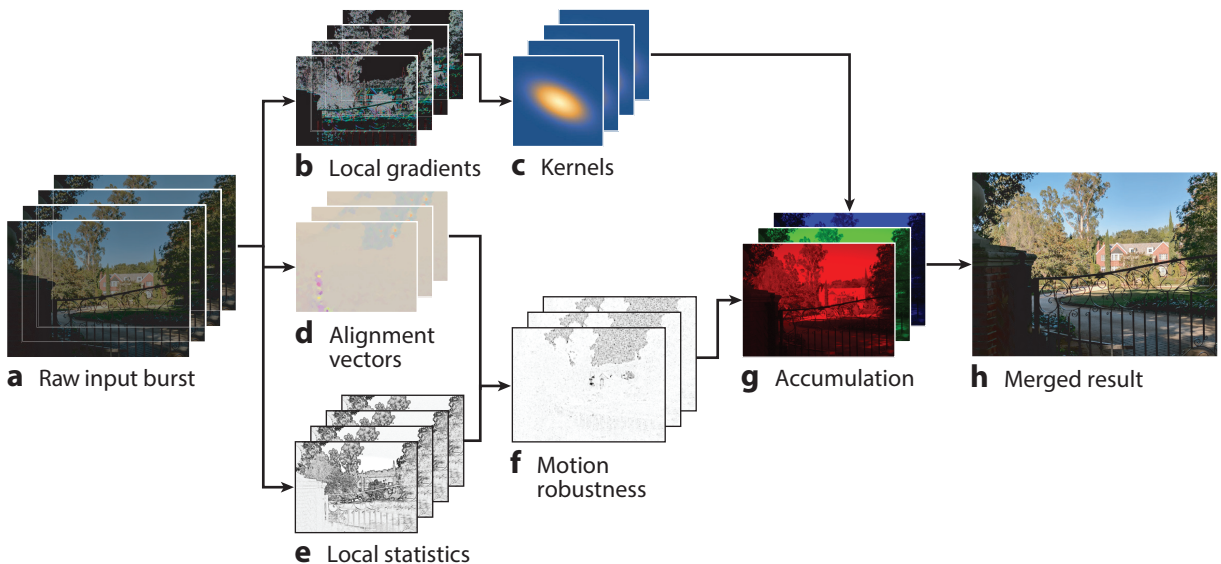


Figure 16

Overview of super-resolution from raw images. (a) A captured burst of raw (Bayer CFA) images is the input to our algorithm. (d) Every frame is aligned locally to a single frame, called the base frame. (g) We estimate each frame's contribution at every pixel through kernel regression. (c) The kernel shapes are adjusted based on (b) the estimated local gradients, (f) and the sample contributions are weighted based on a robustness model. This robustness model computes a per-pixel weight for every frame using (d) the alignment field and (e) local statistics gathered from the neighborhood around each pixel. (b) The final merged RGB image is obtained by normalizing the accumulated results per channel. We call the steps depicted in panels b–g the merge step. Figure adapted from Wronski et al. (2019) (open access).

is aligned locally with a single key frame from the burst (called the base frame). In the super-resolution application, however, the accuracy of the alignment must meet a higher standard. For instance, due to the color subsampling, to super-resolve to even the native sensor grid, the accuracy of the registration must be at worst one-half pixel. This is a significantly higher burden both computationally and statistically, making it difficult, if not impossible, to achieve super-resolution in darker scenes, or at high magnification (beyond $2\times$).

With high-accuracy alignment, each frame's local contributions are estimated through kernel regression (Takeda et al. 2007) and accumulated across the entire burst separately per color plane. To achieve local detail, texture, and geometry recovery, the kernel shapes are adjusted based on the estimated signal features, and the sample contributions are weighted based on a robustness model (which estimates alignment accuracy). Finally, a per-channel normalization yields the merged RGB image.

Super-resolution is arguably not an alien process to the HVS. It would appear that the human brain also processes visual stimuli in a way that allows us to discriminate details beyond the physical resolution given by optics and retinal sampling alone. This is commonly known as visual hyper-acuity, as in the work of Westheimer (1975). A possible mechanism of visual super-resolution is the random eye micromovements known as microsaccades and ocular drifts (Intoy & Rucci 2020, Rucci et al. 2007).

Interestingly, in the super-resolution work described by Wronski & Milanfar (2018), natural hand tremors play a similar role to eye movements. A natural, involuntary hand tremor is always present when we hold any object. This tremor comprises low-amplitude and high-frequency components consisting of a mechanical-reflex component and a second component that causes microcontractions in the limb muscles (Riviere et al. 1998). Wronski et al. (2019) showed that the hand tremor of a user holding a mobile camera is sufficient to provide subpixel movements across the images in a burst for the purpose of super-resolution. Experimental measurements of such tremor in captured bursts of images from a mobile device are illustrated in **Figure 17**.

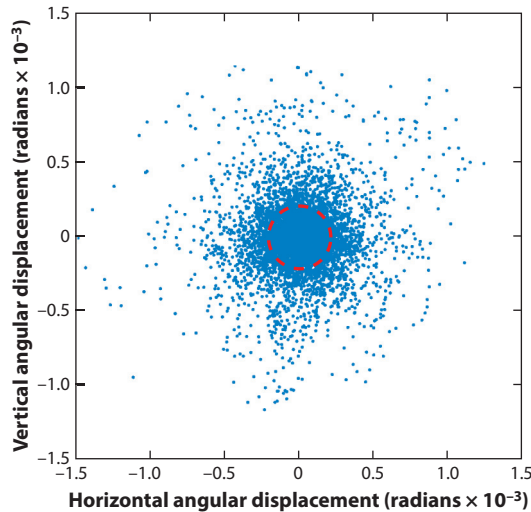


Figure 17

Horizontal and vertical angular displacement (excluding translational displacement) measured from handheld motion across 86 bursts. The red circle corresponds to one standard deviation, or roughly 0.9 pixels. Figure adapted from Wronski et al. (2019) (open access).




Figure 18

Merging a burst onto different target grid resolutions. From left to right: $1\times$, $1.5\times$, $2\times$. The combination of the Super Res Zoom algorithm and the phone's optical system leads to significantly improved results when merged onto a $1.5\times$ grid; small improvements up to $2\times$ zoom are also noted. Figure adapted from Wronski et al. (2019) (open access).

As illustrated in **Figure 18**, the merge algorithm alone is able to deliver resolution comparable to roughly a dedicated telephoto lens at a modest magnification factor, no more than $2\times$. Of course, the same algorithm can also be applied to the telephoto lens itself, again with typically even more modest gains in resolution. This suggests that, to have a general solution for zoom across a broad range of magnifications, a combination of multiframe merge and high-quality single-frame crop-and-upscale is required (see **Figure 15**). This upscaling technology is described below.

4.2.2. Single-frame upscaling. Digital zoom, also known as crop-and-scale, allows the user to change the field of view (FOV) of the captured photograph through digital postprocessing. This operation is essential to allow the photographer to zoom even in cameras that have a dedicated telephoto lens (optical zoom). The operation consists of cropping the image to the desired FOV and then digitally enlarging the cropped region to obtain the desired zoom factor.

A key challenge of digital zoom lies in performing the upscaling (sometimes also called single-image super-resolution) in a way that preserves image details. Traditional techniques are based on the unrealistic assumption that the digital image is generated by sampling a smooth and regular continuous unknown image. This continuous model allows us to generate arbitrary in-between samples from the observed ones by means of interpolation. There are different interpolation schemes that trade off computational cost, quality of the upsampled image (e.g., level of blur introduced), and other possible artifacts. An interpolation scheme is characterized by an interpolation kernel that specifies how the intermediate subpixel sample is computed from the nearby ones. Bilinear interpolation is simple to compute but generates intermediate samples that result in a final image with blurry appearance. At the other extreme, the Lanczos interpolation best approximates the assumed continuous image model but has a higher computational cost since it uses a larger context (large kernel). An advantage of this type of linear interpolation is that intermediate samples can be calculated at arbitrary positions, thus allowing digital zooming of any factor.

An extension of linear interpolation methods that relies on machine learning techniques is rapid and accurate image super resolution (RAISR) (Romano et al. 2017). This method can be seen as a double extension of linear methods. On the one hand, the interpolation kernel is learned from a training data set of pairs of low- and high-resolution images. This is done by finding the

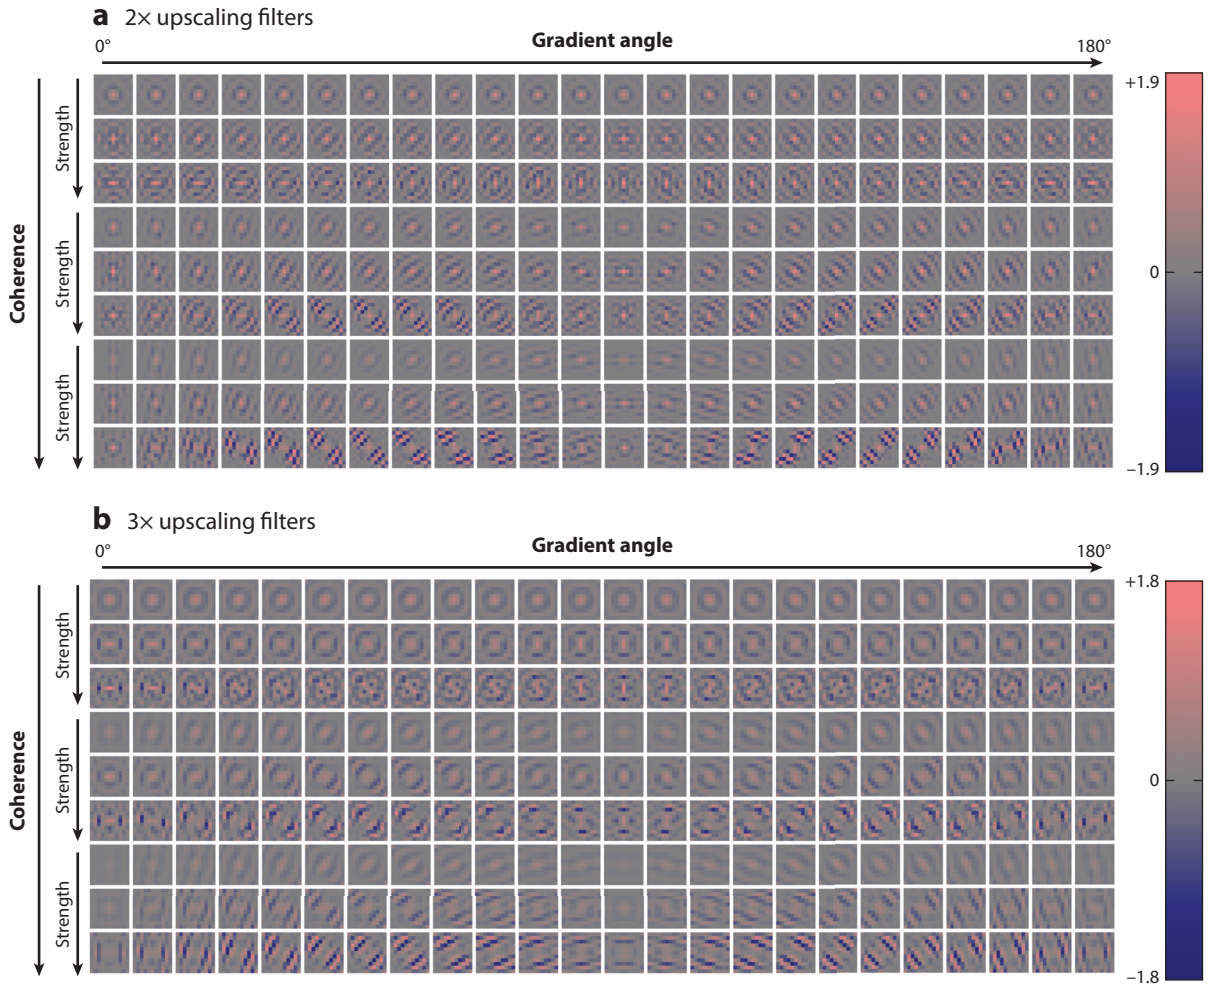


Figure 19

Visualization of the learned filter sets for 2× and 3× upscaling, indexed by angle, strength, and coherence-based hashing of patch gradients. Figure adapted from Romano et al. (2017).

best kernel that minimizes the interpolation error for the image pairs in the training data set. On the other hand, RAISR goes one step further and learns a collection of interpolating kernels, each one specialized for a certain local structure encoded by the gradient strength, orientation, and coherence. Examples of such trained filter banks are shown in **Figure 19**. Per each subset of filters, the angle varies from left to right; the top, middle, and bottom three rows correspond to low, medium, and high coherence. It is important to note that the filters at different upscaling factors are not trivial transformations of one another. For instance, the 3× filters are not derived from the 2× filters—each set of filters carries novel information from the training data. Given the apparent regularity of these filters, it may also be tempting to imagine that they can be parameterized by known filter types (e.g., Gabor). This is not the case. Specifically, the phase response of the trained filter is deeply inherited from the training data, and no parametric form has been found that is able to mimic this generally.

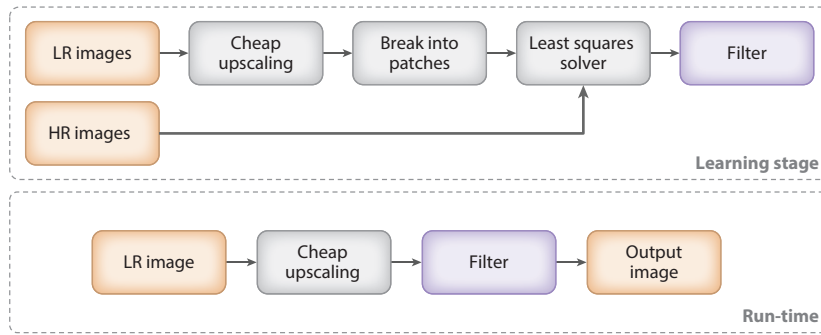


Figure 20

The learning and application of a filter that maps a class of LR patches to their HR versions. More generally, a set of such filters is learned and indexed by local geometric structures, as shown in **Figure 19**. Figure adapted from Romano et al. (2017). Abbreviations: HR, high-resolution; LR, low-resolution.

The overall idea behind RAISR and related methods is that well-defined structures, like edges, can be better interpolated if the interpolation kernel makes use of the specific orientation and local structure properties. During execution, RAISR scans the input image, defines which kernel to use on each pixel, and then computes the upscaled image using the selected kernels on a per-pixel basis. The overall structure of the RAISR algorithm is shown in **Figure 20**.

RAISR is trained to enlarge an input image by an integer factor (2–4×), but it does not allow intermediate zooms. In practice, RAISR is combined with a linear interpolation method (bicubic, Lanczos) to generate the zoom factor desired by the user.

With the advancement of DNNs in recent years, a wide variety of new image upscaling methods have emerged (Wang et al. 2021). Similar to RAISR, deep-learning-based methods propose to learn from image examples how to map a low-resolution image into a high-resolution one. These methods generally produce high-quality results, but they are not as computationally efficient as shallow interpolation methods, such as RAISR, so their use in mobile phones is not yet widespread. Undoubtedly, deep image upscaling is one of the most active areas of research. Recent progress in academic research, combined with more powerful and dedicated hardware, may produce significant improvements that could be part of the next generation of mobile cameras.

4.3. Synthetic Bokeh

One of the main characteristics of mobile phone cameras is that the whole image is either in focus or not. The DoF, defined as the range of depths that are in focus (sharp), is frequently used by photographers to distinguish the main subject from the background. As discussed in Section 2, due to the small and fixed aperture used on smartphone cameras, capturing a shallow-DoF image is virtually impossible.

The range of the DoF is inversely proportional to the size of the camera aperture: A wide aperture produces a shallow DoF, while a narrow aperture leads to a wider DoF. On mobile phone cameras, physical limitations make it impossible to have a wide aperture. This implies that capturing images with a shallow DoF is virtually impossible. Although an all-in-focus image retains the most information, for aesthetic and artistic reasons, users may want to have control of the DoF.

Recently, mobile phone manufacturers have introduced a computational (shallow) DoF effect called a synthetic bokeh (see **Figure 21**). An accurate depth map estimate would enable computationally introducing spatially varying depth blur and simulate the DoF effect. The traditional solution to estimating a depth map is based on stereo vision and requires two cameras. Adding a

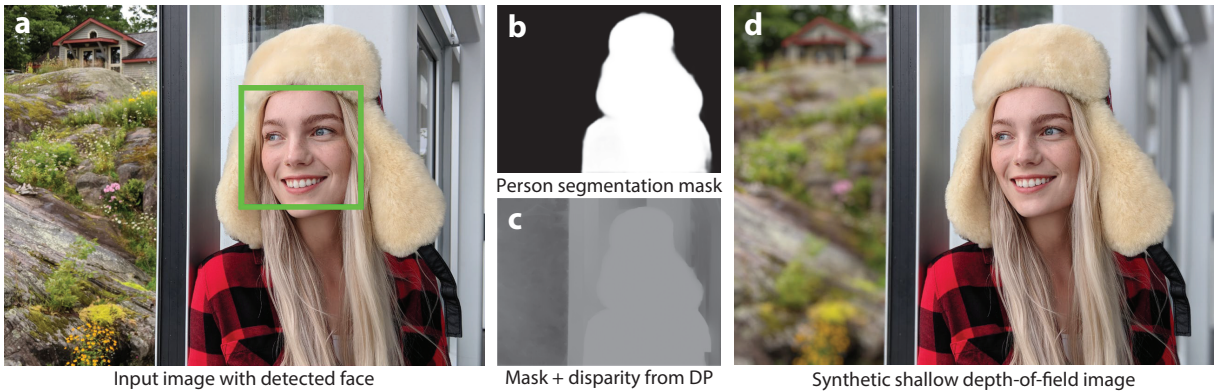


Figure 21

Shallow depth of field can be computationally introduced by blurring (a) an all-in-focus image using (b) segmentation and (c) depth estimation. This effect is known as (d) synthetic bokeh. Figure adapted with permission from Wadhwa et al. (2018).

second camera introduces additional costs and increases power consumption and size. An alternative is to introduce a dedicated depth sensor based on structured light or time-of-flight technologies. However, these tend to be expensive and mainly work indoors, which significantly restricts their use. Accurately estimating a depth map from a single image is a severely ill-posed problem that generally leads to very limited accuracy.

Wadhwa et al. (2018) introduced a system to synthetically generate a DoF effect on smartphone cameras. The method runs completely on the device and uses only the information from a single camera (rear or front facing). The key idea is to incorporate a DNN to segment out people and faces and then use the segmentation to adaptively blur the background. Additionally, if available, the system uses dual-pixel information now present in many hardware auto-focus systems. The dual-pixel data provide very small-baseline stereo information that allows algorithmic generation of dense depth maps.

Meanwhile, the front-facing camera is used almost exclusively to take selfie photos—that is, a close-up of the face and upper half of the photographer’s body. A neural network trained for this type of image allows segmenting the main character out from the background. The background is then appropriately blurred to give the idea of DoF. When using the rear-facing camera, no prior information about the photograph composition can be assumed. Thus, having dense depth information becomes crucial.

It is worth mentioning that this computational DoF system does not necessarily lead to a physically plausible photograph, as would have been taken by a camera with a wider aperture—it merely suggests the right look. For instance, among other algorithm design choices, all pixels belonging to the segmentation mask are assumed to be in focus even if they are at different depths.

5. THE FUTURE AND UPCOMING CHALLENGES

Mobile cameras have made significant strides in quality, matching (and even surpassing) the image quality of so-called micro-4/3 standalone cameras (<https://www.dpreview.com/articles/6476469986/dpreview-products-of-the-year-2018?slide=25>). While it seems unlikely that, absent a major change in form factor, smartphone cameras will equal DSLR quality, much still remains to be done. In this section, we discuss some promising avenues and new challenges.

5.1. Algorithmics

The impressive success of neural networks in computer vision has not (yet) been widely replicated in practical aspects of computational photography. Indeed, the impressive progress in mobile imaging has largely been facilitated by methods that are mostly not based on DNNs. Given the proliferation of DNNs in every other aspect of technology, this may seem surprising. Two observations may help explain this situation. First, DNNs still have relatively heavy computing and memory requirements that are mostly outside the scope of current capabilities of mobile devices. This may change soon. Second, resource limitations aside, DNN-based (particularly the so-called generative) models still have the tendency to produce certain artifacts in the final images that are either undesirable or intolerable in a consumer device. Furthermore, such errors are not easily diagnosed and repaired because, unlike for existing methods, tuning the behavior of deep models is not easy. These issues too will be remedied in due time. Meanwhile, DNN approaches continue to be developed with the aim of replacing the entire end-to-end processing pipeline—examples include DeepISP (Schwartz et al. 2018) and the work of Ignatov et al. (2020), to name just two.

5.2. Curation

Today, nearly everyone who can afford to have a smartphone owns one, and we now take and share more photos than ever. Given how little cost and effort picture taking entails, we have also evolved a tendency to capture multiple pictures of the same subject. Yet typically only a few of our many shots turn out well or to our liking. As such, storage, curation, and retrieval of photographs have become other aspects of photography that have drawn attention recently and deserve much more work. Some recent methods (Talebi & Milanfar 2018) have developed neural network models trained on many images annotated for technical and aesthetic quality, which now enable machine evaluation of images in both qualitative and quantitative terms. Of course, this technology is in very early stages of development and represents aggregate opinion that is not yet necessarily meant to cater to personal taste. Similar models can also rank photos based on their technical quality— aspects such as whether the subject is well lit, centered, and in focus. Needless to say, much work remains to be done in this area.

5.3. Broader Use Cases

The proliferation of cameras and computational photography technology is of course not limited to the mobile platform. It is indeed not an exaggeration to say that cameras are nearly everywhere. Many of the techniques developed for the mobile platform may in fact be useful for enhancing the quality of images derived on other platforms, notably scientific instrumentation, automotive imaging, satellite imaging, and more. However, caution is warranted, and key differences among these platforms should be noted.

In particular, it is important to note that the imaging pipelines developed in the context of mobile photography are specifically tuned for producing aesthetically pleasing images. Meanwhile, in scientific uses of computational photography, the end goal is not the image itself but rather certain, and varied, information extracted from the images. For instance, in the medical realm, the end task may be diagnostic, and this may not be best facilitated by a pretty picture. Instead, what is required is a maximally informative picture. Correspondingly, cameras on mobile devices are not built, configured, or tuned to provide such information.⁵ An interesting case study is the

⁵For instance, AWB can be counterproductive if the end goal is to make a physical measurement that depends on color fidelity in a different sense.

role that cameras have played (or could better have played) in recent environmentally calamitous events such as the massive wildfires in California. Nearly all cameras, tuned to normal viewing conditions and biased toward making the pictures pleasing, were largely unable to capture the true physical attributes of the dark, orange-hued skies polluted with smoke (Beil 2020).

The bottom line is that mobile cameras, properly reimagined and built, can play an even more useful, helpful, and instrumental role in our lives than they do today.

5.4. Epilogue

The technology behind computational photography has advanced rapidly in the past decade—the science and engineering techniques that generate high-quality images from small mobile cameras will continue to evolve. However, so too will our needs and tastes for the types of devices that we are willing to carry around and the kinds of visual or other experiences that we wish to record and share.

It is hard to predict with any certainty what the mobile devices of the future will look like. However, as surely as Ansel Adams would not have seen the mobile phone camera coming, we too may be surprised by both the form and the vast new uses of these devices in the next decade.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

The authors wish to acknowledge the computational imaging community of scholars and colleagues—industrial and academic alike—whose work has led to the advances reported in this review. While we could not cite them all, we dedicate this paper to their collective work.

LITERATURE CITED

- Abdelhamed A, Lin S, Brown MS. 2018. A high-quality denoising dataset for smartphone cameras. In *Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT*, pp. 1692–700. Piscataway, NJ: IEEE
- Adams A. 1935. *Making a Photograph: An Introduction to Photography*. London: Studio Ed. Ltd.
- Afifi M, Price B, Cohen S, Brown MS. 2019. *Image recoloring based on object color distributions*. Work. Pap., York Univ., Toronto, ON, Can.
- Ahn H, Keum B, Kim D, Lee HS. 2013. Adaptive local tone mapping based on Retinex for high dynamic range images. In *Proceedings of the 2013 IEEE International Conference on Consumer Electronics (ICCE), Las Vegas, NV*, pp. 153–56. Piscataway, NJ: IEEE
- Banterle F, Artusi A, Debattista K, Chalmers A. 2017. *Advanced High Dynamic Range Imaging*. Boca Raton, FL: CRC Press
- Barron JT, Tsai YT. 2017. Fast Fourier color constancy. In *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI*, pp. 886–94. Piscataway, NJ: IEEE
- Bayer BE. 1975. *Color imaging array*. US Patent 3,971,065
- Beil K. 2020. Photography has gotten climate change wrong from the start. *The Atlantic*, Nov. 27. <https://www.theatlantic.com/ideas/archive/2020/11/photography-has-never-known-how-handle-climate-change/617224>
- Bruhn A, Weickert J, Schnörr C. 2005. Lucas/Kanade meets Horn/Schunck: combining local and global optic flow methods. *Int. J. Comput. Vis.* 61:211–31
- Buades A, Coll B, Morel JM. 2005. A review of image denoising algorithms, with a new one. *Multiscale Model. Simul.* 4:490–530

- Buades A, Coll B, Morel JM. 2006. *Image enhancement by non-local reverse heat equation*. CMLA Preprint 22, Cent. Math. Leurs Appl., Cachan, France
- Buchsbaum G. 1980. A spatial processor model for object colour perception. *J. Franklin Inst.* 310:1–26
- Burger HC, Schuler CJ, Harmeling S. 2012. Image denoising: Can plain neural networks compete with BM3D? In *Proceedings of the 2012 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2392–99. Piscataway, NJ: IEEE
- Burr D, Ross J, Morrone MC. 1986. Seeing objects in motion. *Proc. R. Soc. Lond. B* 227:249–65
- Cerda X, Parraga CA, Otazu X. 2018. Which tone-mapping operator is the best? A comparative study of perceptual quality. *J. Opt. Soc. Am. A* 35:626–38
- Chatterjee P, Milanfar P. 2010. Is denoising dead? *IEEE Trans. Image Proc.* 19:895–911
- Chen Y, Yu W, Pock T. 2015. On learning optimized reaction diffusion processes for effective image restoration. In *Proceedings of the 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Boston, MA*, pp. 5261–69. Piscataway, NJ: IEEE
- Cheng D, Prasad DK, Brown MS. 2014. Illuminant estimation for color constancy: why spatial-domain methods work and the role of the color distribution. *J. Opt. Soc. Am. A* 31:1049–58
- Cheng D, Price B, Cohen S, Brown MS. 2015. Effective learning-based illuminant estimation using simple features. In *Proceedings of the 2015 IEEE Conference of Computer Vision and Pattern Recognition (CVPR), Boston, MA*, pp. 1000–8. Piscataway, NJ: IEEE
- Dabov K, Foi A, Katkovnik V, Egiazarian K. 2007. Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Trans. Image Proc.* 16:2080–95
- Debevec PE, Malik J. 1997. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH '97: Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 369–78. New York: ACM
- Delbracio M, Garcia-Dorado I, Choi S, Kelly D, Milanfar P. 2020. Polyblur: removing mild blur by polynomial reblurring. arXiv:2012.09322 [cs.CV]
- Efros AA, Leung TK. 1999. Texture synthesis by non-parametric sampling. In *Proceedings of the Seventh IEEE International Conference on Computer Vision, Kerkyra, Greece*, Vol. 2, pp. 1033–38. Piscataway, NJ: IEEE
- El-Henawy I, Amin A, Kareem Ahmed HA. 2014. A comparative study on image deblurring techniques. *Int. J. Adv. Comput. Sci. Technol.* 3:1–8
- Elad M. 2002. On the origin of the bilateral filter and ways to improve it. *IEEE Trans. Image Proc.* 11:1141–50
- Farsiu S, Elad M, Milanfar P. 2006. Multiframe demosaicing and super-resolution of color images. *IEEE Trans. Image Proc.* 15:141–59
- Fergus R, Singh B, Hertzmann A, Roweis ST, Freeman WT. 2006. Removing camera shake from a single photograph. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Papers*, pp. 787–94. New York: ACM
- Gehler PV, Rother C, Blake A, Minka T, Sharp T. 2008. Bayesian color constancy revisited. In *Proceedings of the 2008 IEEE Conference of Computer Vision and Pattern Recognition, Anchorage, AK*, art. 4587765. Piscataway, NJ: IEEE
- Gharbi M, Chaurasia G, Paris S, Durand F. 2016. Deep joint demosaicking and denoising. *ACM Trans. Graph.* 35:191
- Godard C, Matzen K, Uyttendaele M. 2018. Deep burst denoising. In *Proceedings of the 15th European Conference on Computer Vision, Munich, Germany*, pp. 560–77. Berlin: Springer
- Hasinoff SW, Durand F, Freeman WT. 2010. Noise-optimal capture for high dynamic range photography. In *Proceedings of the 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR), San Francisco, CA*, pp. 553–60. Piscataway, NJ: IEEE
- Hasinoff SW, Sharlet D, Geiss R, Adams A, Barron JT, et al. 2016. Burst photography for high dynamic range and low-light imaging on mobile cameras. *ACM Trans. Graph.* 35:192
- Horn BK, Schunck BG. 1981. Determining optical flow. *Artif. Intel.* 17:185–203
- Hosseini MS, Plataniotis KN. 2019. Convolutional deblurring for natural imaging. *IEEE Trans. Image Proc.* 29:250–64
- Hu Y, Wang B, Lin S. 2017. FC⁴: fully convolutional color constancy with confidence-weighted pooling. In *Proceedings of the 2017 IEEE Conference of Computer Vision and Pattern Recognition (CVPR), Honolulu, HI*, pp. 330–39. Piscataway, NJ: IEEE

- Ignatov A, Van Gool L, Timofte R. 2020. Replacing mobile camera ISP with a single deep learning model. In *Proceedings of the 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, Seattle, WA, pp. 2275–85. Piscataway, NJ: IEEE
- Intoy J, Rucci M. 2020. Finely tuned eye movements enhance visual acuity. *Nat. Commun.* 11:795
- Jiang J, Liu D, Gu J, Süsstrunk S. 2013. What is the space of spectral sensitivity functions for digital color cameras? In *Proceedings of the 2013 IEEE Workshop on Applications of Computer Vision (WACV)*, Clearwater Beach, FL, pp. 168–79. Piscataway, NJ: IEEE
- Kainz F, Murthy K. 2019. Astrophotography with night sight on Pixel phones. *Google AI Blog*, Nov. 26. <https://ai.googleblog.com/2019/11/astrophotography-with-night-sight-on.html>
- Karaimer HC, Brown MS. 2016. *A software platform for manipulating the camera imaging pipeline*. Paper presented at the European Conference on Computer Vision (ECCV '16), Oct. 8–16, Amsterdam
- Kelber A, Yovanovich C, Olsson P. 2017. Thresholds and noise limitations of colour vision in dim light. *Philos. Trans. R. Soc. Lond. B* 372:20160065
- Kovácsznay LS, Joseph HM. 1955. Image processing. *Proc. Inst. Radio Eng.* 43:560–70
- Kramer RH, Davenport CM. 2015. Lateral inhibition in the vertebrate retina: the case of the missing neurotransmitter. *PLOS Biol.* 13:e1002322
- Land EH. 1974. The retinex theory of colour vision. *Proc. Roy. Inst. Gr. Br.* 47:23–58
- Lebrun M, Colom M, Buades A, Morel JM. 2012. Secrets of image denoising cuisine. *Acta Numer.* 21:475–576
- Levin A, Weiss Y, Durand F, Freeman WT. 2009. Understanding and evaluating blind deconvolution algorithms. In *Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Miami, FL, pp. 1964–71. Piscataway, NJ: IEEE
- Levoy M, Pritch Y. 2018. Night sight: seeing in the dark on Pixel phones. *Google AI Blog*, Nov. 14. <https://ai.googleblog.com/2018/11/night-sight-seeing-in-dark-on-pixel.html>
- Liba O, Murthy K, Tsai YT, Brooks T, Xue T, et al. 2019. Handheld mobile photography in very low light. *ACM Trans. Graph.* 38:164
- Lindenbaum M, Fischer M, Bruckstein A. 1994. On Gabor's contribution to image enhancement. *Pattern Recognit.* 27:1–8
- Liu D, Wen B, Fan Y, Loy CC, Huang TS. 2018. Non-local recurrent network for image restoration. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NeurIPS 2018)*, Montreal, Canada, pp. 1680–89. New York: ACM
- Liu P, Zhang H, Zhang K, Lin L, Zuo W. 2018. Multi-level wavelet-CNN for image restoration. In *Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, Salt Lake City, UT, pp. 886–95. Piscataway, NJ: IEEE
- Longere P, Zhang X, Delahunt PB, Brainard DH. 2002. Perceptual assessment of demosaicing algorithm performance. *Proc. IEEE* 90:123–32
- Lovejoy B. 2020. What is a periscope lens, and what would it mean for future iPhones? *9to5Mac*, July 22. <https://9to5mac.com/2020/07/22/periscope-lens/>
- Ma K, Yeganeh H, Zeng K, Wang Z. 2015. High dynamic range image compression by optimizing tone mapped image quality index. *IEEE Trans. Image Proc.* 24:3086–97
- Mao X, Shen C, Yang YB. 2016. Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections. In *Proceedings of the 30th International Conference on Neural Information Processing Systems (NIPS 2016)*, Barcelona, Spain, pp. 2810–18. New York: ACM
- Meinhardt T, Moeller M, Hazirbas C, Cremers D. 2017. Learning proximal operators: using denoising networks for regularizing inverse imaging problems. In *Proceedings of the 2017 IEEE International Conference on Computer Vision (ICCV)*, Venice, Italy, pp. 1799–808. Piscataway, NJ: IEEE
- Mertens T, Kautz J, Reeth FV. 2007. Exposure fusion. In *Proceedings of the 15th Pacific Conference on Computer Graphics and Applications (PG'07)*, Maui, HI, pp. 382–90. Piscataway, NJ: IEEE
- Milanfar P. 2013. A tour of modern image filtering: new insights and methods, both practical and theoretical. *IEEE Signal Proc. Mag.* 30:106–28
- Mildenhall B, Barron JT, Chen J, Sharlet D, Ng R, Carroll R. 2018. Burst denoising with kernel prediction networks. In *Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Salt Lake City, UT, pp. 2502–10. Piscataway, NJ: IEEE

- Osher S, Rudin LI. 1990. Feature-oriented image enhancement using shock filters. *SIAM J. Numer. Anal.* 27:919–40
- Paris S, Hasinoff SW, Kautz J. 2011. Local Laplacian filters: edge-aware image processing with a Laplacian pyramid. *ACM Trans. Graph.* 30:68
- Plötz T, Roth S. 2017. Benchmarking denoising algorithms with real photographs. In *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI*, pp. 2750–59. Piscataway, NJ: IEEE
- Poynton C. 2012. *Digital Video and HD: Algorithms and Interfaces*. Amsterdam: Morgan Kaufmann. 2nd ed.
- Ramanath R, Drew MS. 2014. *Von Kries Hypothesis*. Berlin: Springer
- Ratcliff F. 1965. *Mach Bands: Quantitative Studies on Neural Networks*, Vol. 2. San Francisco: Holden-Day
- Reinhard E, Stark M, Shirley P, Ferwerda J. 2002. Photographic tone reproduction for digital images. *ACM Trans. Graph.* 21:267–76
- Remez T, Litany O, Giryes R, Bronstein AM. 2018. Class-aware fully convolutional Gaussian and Poisson denoising. *IEEE Trans. Image Proc.* 27:5707–22
- Riviere C, Rader RS, Thakor NV. 1998. Adaptive cancelling of physiological tremor for improved precision in microsurgery. *IEEE Trans. Biomed. Eng.* 45:839–46
- Romano Y, Isidoro J, Milanfar P. 2017. RAISR: rapid and accurate image super resolution. *IEEE Trans. Comput. Imaging* 3:110–25
- Rucci M, Iovin R, Poletti M, Santini F. 2007. Miniature eye movements enhance fine spatial detail. *Nature* 447:852–55
- Sampat N, Venkataraman S, Yeh T, Kremens RL. 1999. System implications of implementing auto-exposure on consumer digital cameras. In *Proceedings of the SPIE 3650: Sensors, Cameras, and Applications for Digital Photography*, ed. N Sampat, T Yeh, art. 342854. Bellingham, WA: SPIE
- Schwartz E, Giryes R, Bronstein AM. 2018. DeepISP: toward learning an end-to-end image processing pipeline. *IEEE Trans. Image Proc.* 28:912–23
- Smith SM, Brady JM. 1997. SUSAN: a new approach to low level image processing. *Int. J. Comput. Vis.* 23:45–78
- Spira A, Kimmel R, Sochen N. 2007. A short time Beltrami kernel for smoothing images and manifolds. *IEEE Trans. Image Proc.* 16:1628–36
- Stevens SS. 1961. To honor Fechner and repeal his law: A power function, not a log function, describes the operating characteristic of a sensory system. *Science* 133:80–86
- Sun D, Yang X, Liu M, Kautz J. 2018. PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. In *Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Salt Lake City, UT*, pp. 8934–43
- Süsstrunk S, Buckley R, Swen S. 1999. Standard RGB color spaces. In *Proceedings of the Seventh Color Imaging Conference: Color Science, Systems, and Applications*, pp. 127–34. Springfield, VA: Soc. Imaging Sci. Technol.
- Tai Y, Yang J, Liu X, Xu C. 2017. Memnet: a persistent memory network for image restoration. In *Proceedings of the 2017 IEEE International Conference on Computer Vision (ICCV), Venice, Italy*, pp. 4549–57. Piscataway, NJ: IEEE
- Takeda H, Farsiu S, Milanfar P. 2006. Robust kernel regression for restoration and reconstruction of images from sparse, noisy data. In *Proceedings of the 2006 IEEE International Conference on Image Processing, Atlanta, GA*, pp. 1257–60. Piscataway, NJ: IEEE
- Takeda H, Farsiu S, Milanfar P. 2007. Kernel regression for image processing and reconstruction. *IEEE Trans. Image Proc.* 16:349–66
- Talebi H, Milanfar P. 2018. Nima: neural image assessment. *IEEE Trans. Image Proc.* 27:3998–4011
- Tan H, Zeng X, Lai S, Liu Y, Zhang M. 2017. Joint demosaicing and denoising of noisy Bayer images with ADMM. In *Proceedings of the 2017 IEEE International Conference on Image Processing (ICIP), Beijing, China*, pp. 2951–55. Piscataway, NJ: IEEE
- Tellean J, Sullivan A, Yee J, Wang O, Gunawardane P, et al. 2007. Synthetic shutter speed imaging. *Comput. Graph. Forum* 26:591–98
- Tomasi C, Manduchi R. 1998. Bilateral filtering for gray and color images. In *Proceedings of the Sixth IEEE International Conference on Computer Vision, Bombay, India*, pp. 836–46. Piscataway, NJ: IEEE

- Van De Weijer J, Gevers T, Gijsenij A. 2007. Edge-based color constancy. *IEEE Trans. Image Proc.* 16:2207–14
- Wadhwa N, Garg R, Jacobs DE, Feldman BE, Kanazawa N, et al. 2018. Synthetic depth-of-field with a single-camera mobile phone. *ACM Trans. Graph.* 37:64
- Wang Z, Chen J, Hoi SC. 2021. Deep learning for image super-resolution: a survey. *IEEE Trans. Pattern Anal. Mach. Intell.* In press
- Wang Z, Liu D, Yang J, Han W, Huang T. 2015. Deep networks for image super-resolution with sparse prior. In *Proceedings of the 2015 IEEE Conference on Computer Vision (ICCV), Santiago, Chile*, pp. 370–78
- Weickert J. 1999. Coherence-enhancing diffusion. *Int. J. Comput. Vis.* 31:111–27
- Westheimer G. 1975. Visual acuity and hyperacuity. *Investig. Ophthalmol. Vis. Sci.* 14:570–72
- Wronski B, Garcia-Dorado I, Ernst M, Kelly D, Krainin M, et al. 2019. Handheld multi-frame super-resolution. *ACM Trans. Graph.* 38:28
- Wronski B, Milanfar P. 2018. See better and further with super res zoom on the Pixel 3. *Google AI Blog*, Oct. 15. <https://ai.googleblog.com/2018/10/see-better-and-further-with-super-res.html>
- You YL, Xu W, Tannenbaum A, Kaveh M. 1996. Behavioral analysis of anisotropic diffusion in image processing. *IEEE Trans. Image Proc.* 5:1539–53
- Zhang K, Zuo W, Chen Y, Meng D, Zhang L. 2017. Beyond a Gaussian denoiser: residual learning of deep CNN for image denoising. *IEEE Trans. Image Proc.* 26:3142–55
- Zhang K, Zuo W, Zhang L. 2018. FFDNet: toward a fast and flexible solution for CNN-based image denoising. *IEEE Trans. Image Proc.* 27:4608–22
- Zhang Y, Tian Y, Kong Y, Zhong B, Fu Y. 2018. Residual dense network for image restoration. arXiv:1812.10477 [cs.CV]
- Zoran D, Weiss Y. 2011. From learning models of natural image patches to whole image restoration. In *Proceedings of the 2011 IEEE International Conference on Computer Vision (ICCV), Barcelona, Spain*, pp. 479–86. Piscataway, NJ: IEEE