

Handheld Mobile Photography in Very Low Light

ORLY LIBA, KIRAN MURTHY, YUN-TA TSAI, TIM BROOKS, TIANFAN XUE, NIKHIL KARNAD, QIURUI HE, JONATHAN T. BARRON, DILLON SHARLET, RYAN GEISS, SAMUEL W. HASINOFF, Yael PRITCH, and MARC LEVOY, Google Research



Fig. 1. We present a system that uses a novel combination of motion-adaptive burst capture, robust temporal denoising, learning-based white balance, and tone mapping to create high quality photographs in low light on a handheld mobile device. Here we show a comparison of a photograph generated by the burst photography system described in [Hasinoff et al. 2016] and the system described in this paper, running on the same mobile camera. In this low-light setting (about 0.4 lux), the previous system generates an underexposed result (a). Brightening the image (b) reveals significant noise, especially chroma noise, which results in loss of detail and an unpleasantly blotchy appearance. Additionally, the colors of the face appear too orange. Our pipeline (c) produces detailed images by selecting a longer exposure time due to the low scene motion (in this setting, extended from 0.14 s to 0.33 s), robustly aligning and merging a larger number of frames (13 frames instead of 6), and reproducing colors reliably by training a model for predicting the white balance gains specifically in low light. Additionally, we apply local tone mapping that brightens the shadows without over-clipping highlights or sacrificing global contrast.

Taking photographs in low light using a mobile phone is challenging and rarely produces pleasing results. Aside from the physical limits imposed by read noise and photon shot noise, these cameras are typically handheld, have small apertures and sensors, use mass-produced analog electronics that cannot easily be cooled, and are commonly used to photograph subjects that move, like children and pets. In this paper we describe a system for capturing clean, sharp, colorful photographs in light as low as 0.3 lux, where human vision becomes monochromatic and indistinct. To permit handheld photography without flash illumination, we capture, align, and combine multiple frames. Our system employs “motion metering”, which uses an

Authors' address: Orly Liba; Kiran Murthy; Yun-Ta Tsai; Tim Brooks; Tianfan Xue; Nikhil Karnad; Qiurui He; Jonathan T. Barron; Dillon Sharlet; Ryan Geiss; Samuel W. Hasinoff; Yael Pritch; Marc Levoy, Google Research, 1600 Amphitheatre Parkway, Mountain View, CA, 94043.

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

© 2019 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.
0730-0301/2019/11-ART1 \$15.00
<https://doi.org/10.1145/3355089.3356508>

estimate of motion magnitudes (whether due to handshake or moving objects) to identify the number of frames and the per-frame exposure times that together minimize both noise and motion blur in a captured burst. We combine these frames using robust alignment and merging techniques that are specialized for high-noise imagery. To ensure accurate colors in such low light, we employ a learning-based auto white balancing algorithm. To prevent the photographs from looking like they were shot in daylight, we use tone mapping techniques inspired by illusionistic painting: increasing contrast, crushing shadows to black, and surrounding the scene with darkness. All of these processes are performed using the limited computational resources of a mobile device. Our system can be used by novice photographers to produce shareable pictures in a few seconds based on a single shutter press, even in environments so dim that humans cannot see clearly.

CCS Concepts: • **Computing methodologies** → **Computational photography**; **Image processing**;

Additional Key Words and Phrases: computational photography, low-light imaging

ACM Reference Format:

Orly Liba, Kiran Murthy, Yun-Ta Tsai, Tim Brooks, Tianfan Xue, Nikhil Karnad, Qiurui He, Jonathan T. Barron, Dillon Sharlet, Ryan Geiss, Samuel W. Hasinoff, Yael Pritch, and Marc Levoy. 2019. Handheld Mobile Photography in Very Low Light. *ACM Trans. Graph.* 38, 6, Article 1 (November 2019), 16 pages. <https://doi.org/10.1145/3355089.3356508>

1 INTRODUCTION

Low-light conditions present significant challenges for digital photography. This is a well-known problem with a long history of treatment in the camera and mobile device industries, as well as the academic literature.

The fundamental difficulty of low-light photography is that the signal measured by the sensor is low relative to the noise inherent in the measurement process. There are two primary sources of noise in these systems: shot noise and read noise [MacDonald 2006; Nakamura 2016]. If the signal level is very low, read noise dominates. As the signal level rises above the read noise, shot noise becomes the dominant factor. The signal levels relative to these noise sources, that is, the signal-to-noise ratio (SNR), needs to be managed carefully in order to successfully capture images in low light.

The most direct approaches for low-light photography have focused on increasing the SNR of the pixels read from the sensor. These strategies typically fall into one of these categories:

- Expand the camera’s optical aperture.
- Lengthen the exposure time.
- Add lighting, including shutter-synchronized flash.
- Apply denoising algorithms in post-processing.

Unfortunately, these techniques present significant hindrances for photographers—especially for mobile handheld photography, which is the dominant form of photography today. Increasing a camera’s aperture increases the size of the device, making it heavier and more expensive to produce. It also typically increases the length of the optical path, which conflicts with the market’s preference for thin devices. Lengthening the exposure time leads to an increased chance of motion blur, especially when the device is handheld or if there are moving objects in the scene. Adding lighting makes it challenging to produce a natural look: the flash on mobile devices often makes the subject look over-exposed, while the background often looks black because it is too far away for the flash to illuminate. Post-processing single-frame spatial denoising techniques have been widely studied and are often very successful at improving the SNR of an image. However, when noise levels are very high, the reduction of noise often results in significant loss of detail. Moreover, merely improving the SNR addresses only some of the challenges of low light photography and does not yield an overall solution.

This paper presents a system for producing detailed, realistic-looking photographs on handheld mobile devices in light down to about 0.3 lux, a level at which most people would be reluctant to walk around without a flashlight. Our system, which launched to the public in November 2018 as Night Sight on Google Pixel smartphones, is designed to satisfy the following requirements:

- (1) Reliably produce high-quality images when the device is handheld, avoiding visual artifacts.
- (2) Avoid using an LED or xenon flash.
- (3) Allow the user to capture the scene without needing any manual controls.
- (4) Support total scene acquisition times of up to ~6 seconds.
- (5) Accumulate as much light as possible in that time without introducing motion blur from handshake or moving objects.
- (6) Minimize noise and maximize detail, texture, and sharpness.

- (7) Render the image with realistic tones and white balance despite the mismatch between camera and human capabilities.
- (8) Operate on a mobile device (with limited compute resources), with processing time limited to ~2 seconds.

While previous works have focused on specific aspects of low-light imaging (such as denoising [Dabov et al. 2007; Mildenhall et al. 2018]), there have been relatively few works that describe photographic *systems* that address more than one of the above requirements. The main strategies that have been used by previous systems are burst imaging [Hasinoff et al. 2016], and end-to-end trained convolutional neural networks (CNNs) [Chen et al. 2018]. Burst imaging systems capture, align, and merge multiple frames to generate a temporally-denoised image. This image is then further processed using a series of operators including additional spatial denoising, white balancing, and tone mapping. CNN-based systems attempt to perform as much of this pipeline as possible using a deep neural network with a single raw frame as its input. The network learns how to perform all of the image adjustment operators. CNN-based solutions often require significant computational resources (memory and time) and it is challenging to optimize their performance so that they could run on mobile devices. Additionally, end-to-end learning based systems don’t provide a way to design and tune an imaging pipeline; they can only imitate the pipeline they were trained on.

Our solution builds upon the burst imaging system described in [Hasinoff et al. 2016]. In contrast to CNN-based systems, [Hasinoff et al. 2016] was designed to perform well on mobile devices. Simply extending it to collect more frames allows it to generate high-quality photographs down to about 3 lux. However, to produce high-quality images in even lower light, tradeoffs between noise and motion blur must be addressed. Specifically, noise can be improved by increasing the exposure time, but individual frames might contain blur, and because the cited system does not perform deblurring, the motion blur present in the reference frame will remain in the output. Alternatively, noise can be improved by merging more frames and increasing the contribution of each merged frame, effectively increasing temporal denoising, however, doing so will cause moving objects to appear blurry in the output image.

In addition to the problem of low SNR, low-light photographs suffer from a natural mismatch between the capabilities of a digital camera and those of the human visual system. Low light scenes often contain strongly colored light sources. The brain’s visual processing can adapt to this and still pick out white objects in these conditions. However, this is challenging for digital cameras. In darker environments (such as outdoors under moonlight), a camera can take a clean, sharp, colorful pictures. However, we humans don’t see well in these conditions: we stop seeing in color because the cone cells in our retinas do not function well in the dark, leaving only the rod cells that cannot distinguish between different wavelengths of light [Stockman and Sharpe 2006]. This scotopic vision system also has low spatial acuity, which is why things seem indistinct at night. Thus, a long-exposure photograph taken at night will look to us like daytime, an effect that most photographers (or at least most mobile camera users) would like to avoid.

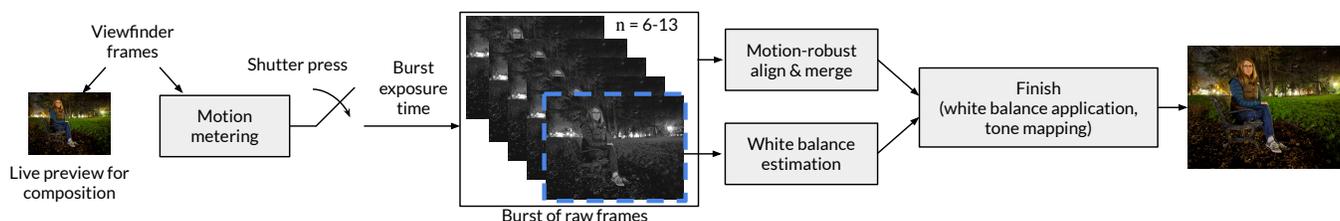


Fig. 2. An overview of our processing pipeline, showing how we extend [Hasinoff et al. 2016]. Viewfinder frames are used for live preview for composition and for motion metering, which determines a per-frame exposure time that provides a good noise vs. motion blur tradeoff in the final result (Section 2). Based on this exposure time, we capture and merge a burst of 6–13 frames (Section 3). The reference frame is down-sampled for the computation of the white-balance gains (Section 4). White balance and tone mapping (Section 5) are applied at the “Finish” stage, as well as demosaicing, spatial and chroma denoising and sharpening. Total capture time is between 1 and 6 seconds after the shutter press and processing time is under 2 seconds.

The system described here solves these problems by extending [Hasinoff et al. 2016] in the following ways:

- (1) **Real-time motion assessment (“motion metering”)** to choose burst capture settings that best balance noise against motion blur in the final result. As an additional optimization, we set limits on exposure time based on an analysis of the camera’s physical stability using on-device gyroscope measurements.
- (2) **Improved motion robustness during frame merging** to alleviate motion blur during temporal denoising. Contributions from merged frames are increased in static regions of the scene and reduced in regions of potential motion blur.
- (3) **White balancing for low light** to provide accurate white balance in strongly-tinted lighting conditions. A learning-based color constancy algorithm is trained and evaluated using a novel metric.
- (4) **Specialized tone mapping** designed to allow viewers of the photograph to see detail they could not have seen with their eyes, but to still know that the photograph conveys a dark scene.

The collection of these new features, shown in the system diagram in Figure 2, enables the production of low-light photographs that exceed human vision capabilities. The following sections cover these topics in detail. Section 2 describes the motion metering system, and Section 3 describes the motion-robust temporal merging algorithm. Section 4 details how we adapt a learning-based color constancy algorithm to perform white balancing in low-light scenes. Section 5 presents our tone-mapping strategy. Section 6, and the supplementary material, show the results of our system for different scenes and acquisition conditions as well as comparisons to other systems and devices. All images were captured on commercial mobile devices unless otherwise specified. Most of the images in the paper were captured while handheld. If a photo was captured using a tripod or braced, it is indicated as “camera stabilized” in the caption.

To further the exploration of low-light photography, we have released a dataset of several hundred raw bursts of a variety of scenes in different conditions captured and processed with our pipeline [Google LLC 2019].

2 MOTION METERING

To capture a burst of images on fixed aperture mobile cameras, the exposure time and gain (ISO) must be selected for each frame in the burst. Here, the gain represents the combination of analog and digital gain. However, several constraints limit the degrees of freedom within these parameters:

- We leverage the strategy in [Hasinoff et al. 2016], where all frames in the burst are captured with the same exposure time and gain after the shutter press, making it easier to align and merge the burst.
- An auto-exposure system (e.g., [Schulz et al. 2007]) determines the sensor’s target sensitivity (exposure time \times gain) according to the scene brightness.
- As mentioned in Section 1, the total capture time should be limited (≤ 6 seconds for our system).

Thus, selecting the capture settings reduces to decomposing the target sensitivity into the exposure time and gain, and then calculating the number of frames according to the capture time limit (i.e. 6 seconds divided by the exposure time, limited to a maximum number of frames set by the device’s memory constraints). [Hasinoff et al. 2016] performs the decomposition using a fixed “exposure schedule”, which conservatively keeps the exposure time low to limit motion blur. However, this may not yield the best tradeoff in all scenarios.

In this section, we present “motion metering”, which selects the exposure time and gain based on a prediction of future scene and camera motion. The prediction is used to shape an exposure schedule dynamically, selecting slower exposures for still scenes and faster exposures for scenes with motion. The prediction is driven by measurements from a new motion estimation algorithm that directly computes the magnitude of the optical flow field. Only the magnitude is needed since the direction of motion blur has no effect on final image quality.

A key challenge is measuring motion efficiently and accurately, as the motion profile in the scene can change rapidly. This can occur due to hand shake or sudden changes in the motion of subjects. Recent work in optical flow has produced highly accurate results using convolutional neural networks ([Dosovitskiy et al. 2015; Revaud et al. 2015]), but these techniques have a high computational cost that prohibits real-time performance on mobile devices. Earlier

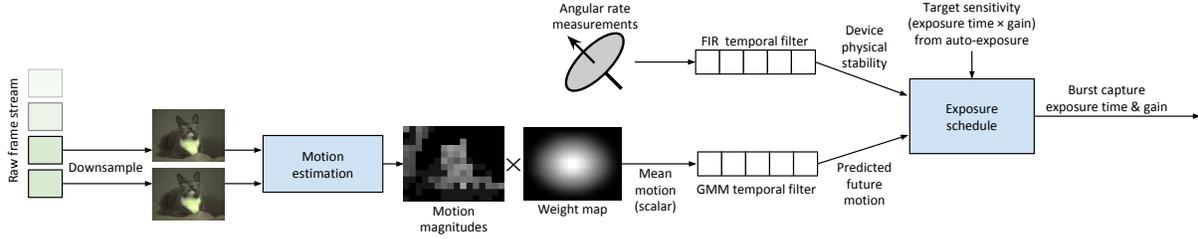


Fig. 3. Overview of the motion metering subsystem. The goal is to select an exposure time and gain (ISO) that provides the best SNR vs. motion blur tradeoff for a burst of frames captured after the shutter press. This is done by measuring scene and camera motion before the shutter press, predicting future motion, and then using the prediction to calculate the capture settings. To measure motion, a stream of frames is downsampled and then processed by an algorithm that directly estimates the magnitude of the optical flow field (“Bounded Flow”). A weighted average of the flow magnitude is then calculated, assigning higher weights to areas where limiting motion blur is more important. To predict future motion, the sequence of weighted averages is processed by a temporal filter based on a Gaussian mixture model (GMM). Concurrently, angular rate measurements from a gyro sensor are filtered to assess the device’s physical stability. These measurements are used to build a custom “exposure schedule” that factorizes the sensor’s target sensitivity computed by a separate auto-exposure algorithm into exposure time and gain. The schedule embodies the chosen tradeoff between SNR and motion blur based on the motion measurements.

optical flow algorithms such as [Lucas and Kanade 1981] are faster but less accurate.

Another key challenge is how to leverage motion measurements to select an exposure time and gain. [Portz et al. 2011] optimizes the exposure time to limit motion blur exclusively at the expense of noise, and [Gryaditskaya et al. 2015] uses a heuristic decision tree driven by the area taken up by motion in the scene. [Boracchi and Foi 2012] provides a technique to optimize the exposure time given a camera motion trajectory and known deblurring reconstruction error in different situations. However, this is primarily useful when strong deblurring can be applied without producing artifacts (a challenging open problem).

Figure 3 depicts our subsystem that addresses these challenges. It combines efficient and accurate motion estimation, future motion prediction, and a dynamic exposure schedule that allows us to tune the system to prioritize noise vs. motion blur. Flow magnitudes are estimated between N successive pairs of frames before shutter press (Section 2.1). Each flow magnitude solution is center-weighted averaged, and the sequence of averages is used to fit a probabilistic model of the motion profile. The model is then used to calculate an upper bound of the minimum motion over the next K frames, where K is chosen such that the motion blur of the merged burst is below some threshold (Section 2.2). Finally, this bound is used to shape the exposure schedule (Section 2.4). We also utilize angular rate estimates from gyro measurements to enhance the future motion prediction (Section 2.3), which improves the robustness of our system in very low light conditions.

2.1 Efficient motion magnitude estimation

We present a new algorithm that directly computes the magnitude of the optical flow field. We start with the Taylor expansion of the brightness constancy constraint [Negahdaripour and Yu 1993]:

$$\Delta I_t(x, y) = \vec{g}(x, y)^T \vec{v}(x, y), \quad (1)$$

the change in intensity between a pair of images at the same pixel location ΔI_t is equal to the inner product of the spatial intensity gradient at that location \vec{g} and the motion vector at that location \vec{v} . The Cauchy-Schwartz inequality can be applied to this expression,

and then rearranged to bound the magnitude of the motion vector:

$$\begin{aligned} |\Delta I_t(x, y)| &\leq \|\vec{g}(x, y)\| \cdot \|\vec{v}(x, y)\| \\ \|\vec{v}(x, y)\| &\geq \frac{|\Delta I_t(x, y)|}{\|\vec{g}(x, y)\|} \end{aligned} \quad (2)$$

Thus, a lower bound for the flow magnitude can be computed by dividing the magnitude of the change in intensity by the norm of the spatial gradient. We call this technique “Bounded Flow”. The bound reaches equivalence when the direction of motion is parallel to the gradient. This is the same property as the aperture problem in full flow solutions: flow along edges can be under-estimated if the direction of motion has a component perpendicular to the gradient. Thus, the “lower bound” in Equation 2 can be considered a direct computation of the optical flow magnitude according to Lucas-Kanade.

The flow magnitude is computed using a pair of linear (non-tone mapped) images on a per-pixel basis. The linearity of the signal enables the efficient modeling of pixel noise variance, where the variance is a linear function of the signal level [Nakamura 2016]. We leverage this by masking out motion in regions where the gradient $\|\vec{g}\|$ is too low relative to the noise standard deviation σ :

$$\|\vec{g}(x, y)\| < K\sigma. \quad (3)$$

We use $K = 2.5$. Discarding small spatial gradients is particularly important in low-light scenes, where noise can be prevalent and can corrupt the flow magnitude estimate. As a final step, the motion magnitudes are refined by downsampling and taking the 90-th percentile motion in each bin, as a form of outlier rejection. Figure 4 shows the effect of applying the noise mask and the refinement.

Table 1 shows the accuracy of the motion magnitudes computed by Bounded Flow relative to other methods, computed from the MPI-Sintel optical flow dataset [Butler et al. 2012]. The methods’ run times are also provided. Bounded Flow with noise masking compares favorably to other flow methods when comparing motion magnitudes estimated from noisy imagery. Bounded Flow is more accurate than other methods that do not use CNNs, but is also faster by at least $2\times$ over Lucas-Kanade based methods.

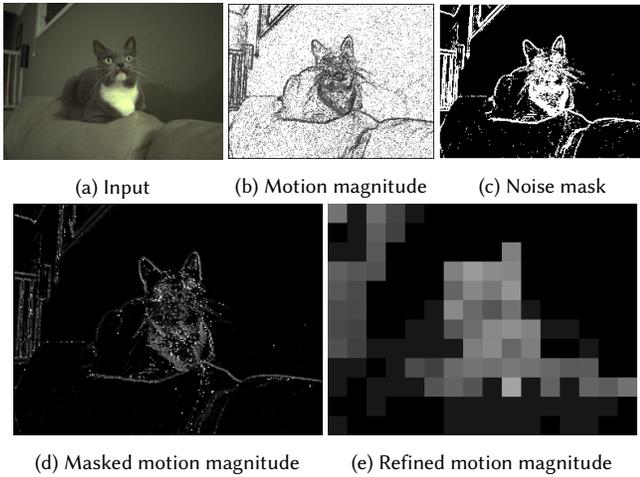


Fig. 4. Intermediate results of “motion metering”. The input is a pair of down-sampled linear images with a known noise variance (a). The white/black levels and the relative exposures are used to normalize the images so that Equation 2 can be applied. The result is (b), which shows the motion magnitude in a grayscale image (white = 10 pixels of motion). The motion is overestimated in areas with low gradients. The noise model is used to compute areas where the gradient is high relative to the noise, and this serves as a mask (c) where valid motion measurements can be made (Equation 3). The motion magnitudes after masking are shown in (d). These are further refined by downsampling with outlier rejection, taking the 90th-percentile motion in each bin. The final result is shown in (e). The resolution of the final output is 16×12 .

Table 1. Comparison of different optical flow algorithms. The means of absolute errors (MAEs) between the predicted and true motion magnitudes were generated using the MPI-Sintel optical flow dataset [Butler et al. 2012]. The images were inverse-tone mapped into a simulated linear space using [Brooks et al. 2019], and high-magnitude noise was added, with a standard deviation of 5% of the white level. The images were then downsampled to 256×192 using a triangle kernel before being processed by the algorithms. Errors are computed only at locations indicated by the noise-based mask in Equation 3. Run times are also provided, measured from single-threaded executions running on an Intel Xeon E5-2690 v4 CPU. We see that Bounded Flow is as accurate as other methods that are not based on neural networks. Furthermore, Bounded Flow is $2 \times$ faster than the next fastest method. This makes it suitable for our real-time application where the motion magnitude needs to be monitored constantly on a mobile device while drawing minimal power.

Motion estimation method	MAE (pixels)	Run time (ms)
Coarse-to-fine flow [Liu 2009]	0.74	310.0
FlowNet [Dosovitskiy et al. 2015]	0.55	451.0
Non-iterative Lucas-Kanade (LK)	0.89	6.25
Iterative LK (max 20 iter)	0.73	9.45
Iterative pyramid LK (max 20 iter, 3 lvls)	0.81	17.0
Our Bounded Flow algorithm	0.62	3.11

After motion magnitude has been measured across the scene, the measurements are aggregated into a single scalar value by taking a weighted average. Typically, a center-weighted average is used, where motion in the center of the scene is weighted highly, and motion on the outside of the scene is discounted. Humans viewing photographs are particularly perceptive to motion blur on faces, therefore, if a face is detected in the scene by a separate automatic system such as [Howard et al. 2017], then the area of the face is weighted higher compared to the rest of the scene. The same type of weighting occurs if a user manually selects an area of the scene. Spatial weight maps of this kind are common in many auto-exposure and autofocus systems.

2.2 Future motion prediction

Given a sequence of weighted average motion measurements, we predict the scene motion of future frames comprising the post-shutter burst. The predictor considers a key property of temporal denoising systems such as [Hasinoff et al. 2016]: one frame out of the first K frames of the burst is selected as the “reference frame”, and is denoised by merging into it other frames in the burst. The end result is that the motion blur of the merged frame is close to the motion blur of the reference frame, and thus, the frame with the least motion blur should be selected as the reference frame. Thus, we can pose the future motion prediction problem as the following: given motion samples v_i , $i = 1, 2, \dots, N$ before the shutter press, what is the minimum motion v_{\min} within the next K frames.

One way to predict future motion is to use a finite impulse response (FIR) filter on the past N samples. There are two issues with this approach. First, FIR filters are not very robust to outliers. If one of the past frames has extremely large motion, either due to a fast moving object or inaccurate motion estimation, the FIR filter output will suddenly increase. However, the outlier frame can easily be rejected when selecting the reference frame, so this should not suddenly lower the final selected exposure time. Second, FIR filters do not consider that only one out of K frames are used as a reference and they require tuning when K changes. For example, when K increases, the maximum exposure time should increase, since there is a higher chance that one frame out of the larger pool of frames remains sharp. However, re-tuning an FIR filter to exhibit this behavior is not straightforward.

To resolve these issues, we propose a motion prediction method based on a cumulative distribution function (CDF). We find the minimum motion v_{\min} such that:

$$Pr \left[v_{\min} \geq \min_{k=1}^K v_k \mid \{v_i\}_{i=1, \dots, N} \right] \geq P_{\text{conf}}. \quad (4)$$

That is, v_{\min} is an upper bound on the minimum motion over the next K frames, to a high degree of confidence (P_{conf}). An upper bound was chosen in order to obtain a conservative estimate of how much the exposure time can be lengthened, thus reducing the chance of unwanted motion blur.

To model the probability of scene motion after the shutter press, we assume that future motion shares the same distribution as scene motion before shutter press. We fit a three-cluster Gaussian mixture model (GMM) to the weighted average motion samples before shutter press using expectation-maximization [Dempster et al. 1977].

Table 2. Comparison of three different filters for future motion prediction, evaluated on a dataset of 777 sequences of weighted average motion magnitude estimates from Bounded Flow taken from sequential frames when composing real-world scenes. For each sequence, we use the motion estimates between N successive frames to predict the minimum pairwise motion of next K frames (here, $N = 10$ and $K = 4$). We compare: 1) the averaging of 5 motion samples, 2) a 8-tap Hamming-window FIR filter with a cutoff frequency of 0.01 of Nyquist, and 3) our Gaussian mixture model (GMM)-based probabilistic prediction on two evaluation metrics: RMSE (root-mean-square error) and RMSE on log2 motion. We calculate log2 to roughly denote f-stops, as the ideal exposure time is inversely proportional to motion magnitude Equation 6. On both metrics, the GMM filter has the lowest error among all three algorithms.

	mean filter	FIR filter	GMM
RMSE	1.2652	1.4671	1.1722
RMSE of log2 motion	0.2538	0.2967	0.1963

Then, assuming that each motion magnitude v_k is independently drawn from the GMM, we have:

$$1 - Pr[v_{\min} \leq v_k]^K \geq P_{\text{conf}}. \quad (5)$$

Since v is a 1D GMM, we can calculate $Pr[v_{\min} \leq v_k]$ for each Gaussian cluster by using the unit normal table. To find the minimum motion v_{\min} such that Equation 5 still holds, we use binary search because the left hand side of Equation 5 is a monotonic function with respect to v_{\min} . Table 2 shows that this GMM-based filter is more accurate at future motion prediction than FIR filtering.

2.3 Camera stability detection

In low light, experienced photographers will often stabilize their camera on a tripod or against a surface to allow for longer exposure times. While we prioritize handheld capabilities, we also allow longer exposures if a user chooses to brace the device, leave it propped, or put the device on a tripod. When the device is handheld, we capture up to 333 ms exposures; when we detect the device is stabilized, we capture up to 1 s exposures. Figure 5 shows a very dark scene captured using a tripod at different exposure times.

Before extending the exposure time, we must be certain that the camera is braced or on a tripod. To accomplish this, we leverage measurements from the device’s gyroscope sensor. Gyroscopes provide a useful signal in video stabilization [Karpenko et al. 2011; Rawat and Jyoti 2011] and alignment [Ringaby and Forssén 2014], and they are similarly valuable for detecting camera stability. Moreover, they provide reliable measurements even in very low light, in which longer exposures have the most benefit.

To assess camera stability from gyro measurements, we temporally smooth angular speeds by averaging them over the $t_0 = 1.466$ s before shutter press. When on a tripod, significant vibration can be caused by the shutter button press, which, on our device, we determined empirically to last for a few hundred milliseconds. This vibration is particularly large if the camera is on a small, lightweight tripod, and can even exceed the camera motion when handheld. Since there is some latency between the physical shutter press and the software receiving this signal, some of these measurements are counted as pre-shutter angular speed. The combination of these

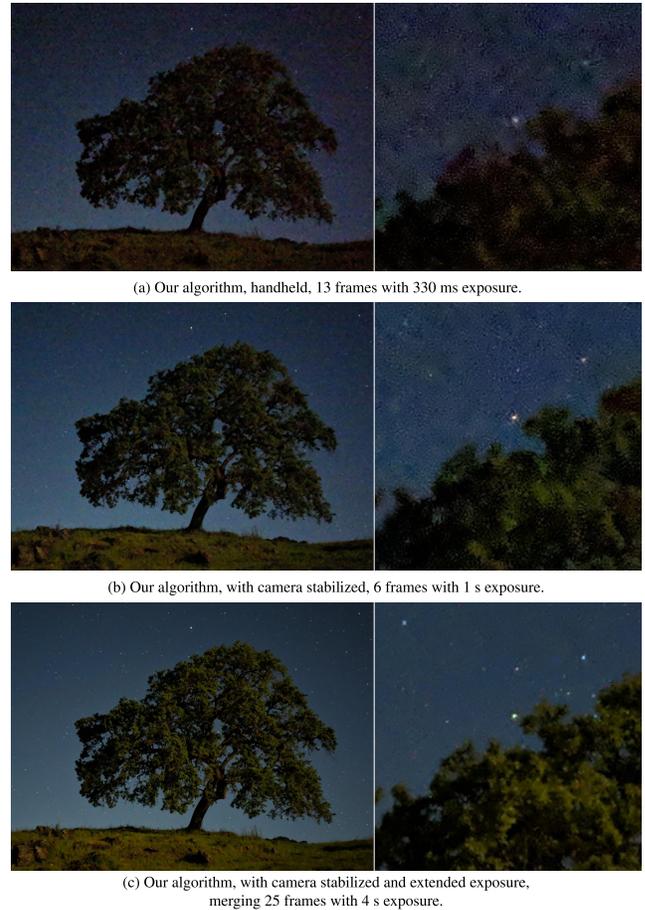


Fig. 5. A demonstration of our algorithm with extended exposures enabled by camera stability detection. The measured light levels at the tree were 0.06 lux. (a) A handheld capture uses 13×333 ms exposures. (b) When stabilized on a tripod, our algorithm uses 6×1 s exposures. The longer individual exposures greatly reduce noise. (c) Camera stability detection could be used for even longer captures, such as the 25×4 s exposures in this example. Noise is significantly reduced, at the cost of a much longer total capture time.

factors makes it difficult to reliably differentiate between handheld and stabilized cameras. We therefore mask out the $t_1 = 0.400$ s of gyroscope readings nearest to the shutter press signal.

The camera motion measurement is thus the average angular speed in the range $-t_0 \leq t \leq -t_1$, where $t = 0$ represents the time at which capture begins. Figure 6 shows the improved reliability of gyroscope measurements after temporal smoothing and shutter press masking. Any residual motion in the first frame of the burst caused by the shutter press will be rejected by reference frame selection.

2.4 Exposure time selection

Based on motion prediction and camera stability detection, our goal is to select the exposure time and gain that will best trade off SNR

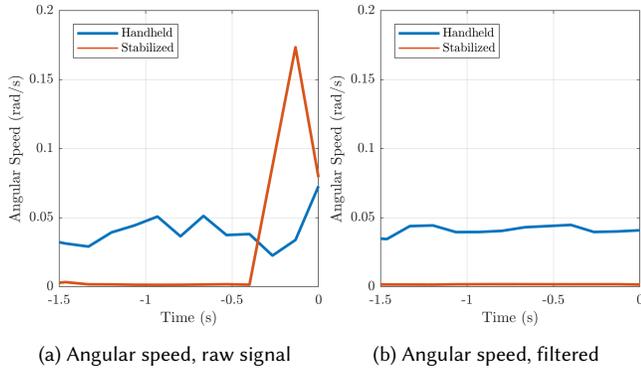


Fig. 6. Masking and smoothing of a gyroscope signal for a more reliable camera stability detection. (a) shows the raw gyroscope signals leading up to the capture of handheld and stabilized bursts taken of the same scene. Despite the camera being stabilized on a tripod, angular speed peaks right before capture due to the shutter press. As described in Section 2.3, motion caused by the shutter press can briefly meet or exceed handheld motion, making reliable camera stability detection challenging. (b) shows the temporally smoothed and masked gyroscope signals, to ignore the brief motion from shutter press. This results in more accurate and reliable discrimination between handheld and stabilized cameras.

for motion blur in the final rendered image. This becomes more challenging in low-light scenes with fast-moving subjects. Selecting a short exposure time that freezes the motion would cause excessive noise, and selecting long exposure times would cause too much motion blur. Therefore, in these scenarios, a balance needs to be struck between the two extremes.

To address this, we use an exposure schedule that factorizes the target sensitivity (generated by an auto-exposure algorithm) into exposure time and gain. Figure 7 depicts a typical exposure schedule for a mobile camera. Instead of using a static exposure schedule, we use a dynamic schedule that is shaped according to the scene motion and camera stability estimates, shown in Figure 7. The camera stability estimate governs the maximum exposure time returned by the schedule. Then, the motion prediction is used to compute a “blur-limiting exposure time”, which restricts the estimated motion blur to B pixels:

$$t_{\text{exp}}^{(B)} = \frac{B}{v_{\text{min}}}, \quad (6)$$

where v_{min} is calculated from Equation 5. As the scene becomes darker, higher gains are required to restrict the exposure time to this value. At a certain point, the gain starts to result in excessive noise, and thus, the exposure schedule begins to increase both exposure time and gain together. Although this can cause the captured images to exceed the motion blur threshold, the noise remains reasonable in the final result.

The difference between static and dynamic schedules depends on the light level. In bright scenes (> 100 lux), the SNR can be high even with short exposure times, so the gain is kept at the minimum value in both schedules. In very dark scenes (< 0.3 lux), both schedules choose the maximum exposure time. The lux levels in between are where the dynamic schedule helps the most.

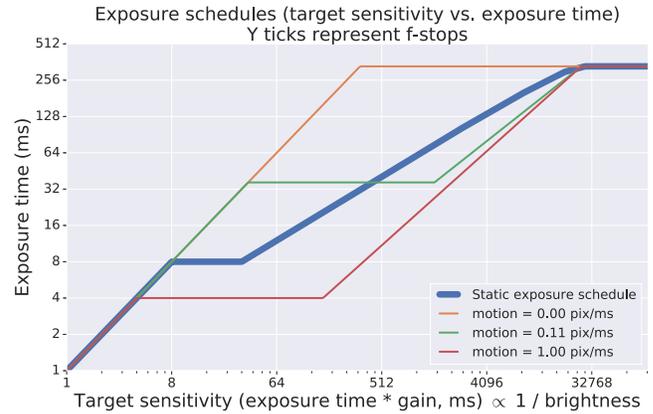


Fig. 7. Cameras typically encode the exposure time vs. gain tradeoff with a static “exposure schedule”, depicted by the thicker blue line. If the exposure schedule curve is higher, the camera favors higher SNR and more motion blur, and if the curve is lower, the camera favors lower SNR and less motion blur. Instead of a static schedule, we use a dynamic schedule based on the measured scene motion. This allows the schedule to provide a better tradeoff for the scene’s conditions. Zero motion (orange line) leads to the slowest possible schedule, and as motion increases, the schedule favors proportionally faster exposure times. The flat regions of the dynamic exposure schedules denote the “blur-limiting exposure time” in Equation 6.

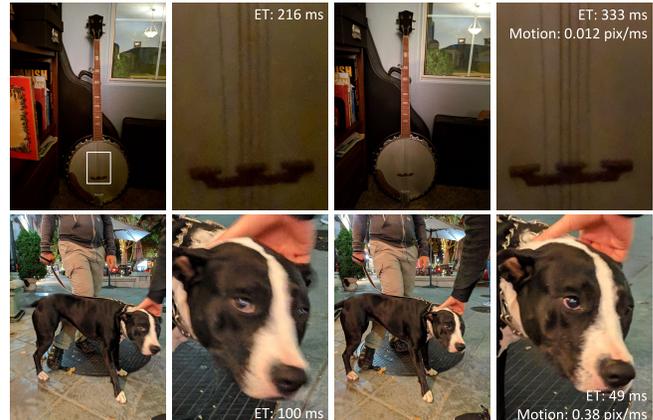


Fig. 8. Comparison of images captured with static (left) and dynamic (right) exposure schedules. The pairs were captured at the same time with different exposure times governed by the schedules. The top row shows the comparison on a still scene (although there is some hand motion) where the dynamic schedule selects an exposure time 0.5 f/stop longer than the static schedule, thereby reducing noise and increasing detail without incurring additional motion blur. The bottom row shows a comparison on a scene with significant subject motion, for which the dynamic schedule selects an exposure time 1 f/stop shorter, thereby reducing motion blur without adding noticeable noise.

Figure 8 shows images collected at the same time with static and dynamic exposure schedules. They indicate that the dynamic schedule chooses better points in the SNR vs. motion blur tradeoff than the static schedule.

3 MOTION-ADAPTIVE BURST MERGING

Aligning and merging a burst of frames is an effective method for reducing noise in the final output. A relatively simple method that aligns and averages frames was implemented in SynthCam [Levoy 2012]. While this method is very effective at reducing noise, moving subjects or alignment errors will cause the resulting image to have motion artifacts. Convolutional neural networks [Aittala and Durand 2018; Godard et al. 2018; Mildenhall et al. 2018] are a more recent and effective method for motion-robust frame merging. However, these techniques require considerable computation time and memory, and are therefore not suited to running within 2 seconds on a mobile device. Bursts can be aligned and merged for other applications, such as super-resolution [Farsiu et al. 2004; Irani and Peleg 1991; Wronski et al. 2019]. These techniques can also be beneficial for denoising in low light and we have provided additional information on using [Wronski et al. 2019] in the supplementary materials.

Since our method must be fast and robust to motion artifacts, we build upon the Fourier domain temporal merging technique described in [Hasinoff et al. 2016], which was targeted for mobile devices. Other fast burst denoising techniques include [Liu et al. 2014], [Delbracio and Sapiro 2015], [Dabov et al. 2007], and [Maggioni et al. 2012]. However, these methods operate on tone-mapped (JPEG) images instead of raw images. As a result, [Hasinoff et al. 2016] benefits from increased dynamic range and better adaptation to different SNR levels in the scene due to having a more accurate noise variance model. In addition, [Liu et al. 2014] performs per-pixel median selection instead of denoising a reference frame, which can lead to artifacts at motion boundaries and locations where different objects with similar intensities occur. For an in-depth comparison of these fast burst denoising techniques, the reader is referred to the supplementary material from [Hasinoff et al. 2016].

Fourier domain merging has a fundamental trade off between denoising strength and motion artifacts. That is, if the scene contains motion, it is challenging to align different frames correctly, and a simple average of the aligned frames would result in motion blur and ghosting. [Hasinoff et al. 2016] addresses this by reducing the degree of merging (i.e., reducing weights in a weighted average) from frequency bands whose difference from the reference frame cannot be explained by noise variance. However, the reduction of the weights is controlled by a global tuning parameter, and as a result, must be tuned conservatively to avoid motion artifacts. Consequently, more denoising could occur in static regions of scenes without incurring artifacts.

In low-light scenes, it becomes important to recover the lost denoising performance in static regions of scenes while remaining robust to motion artifacts. To do so, we propose adding a spatial domain similarity calculation (in the form of “mismatch maps”, defined below) before frequency domain merging. The mismatch maps allow us to specify where the frequency-based merging can be more or less aggressive, on a per-tile basis. Using this method we can increase temporal merging in static regions in order to reduce noise, while avoiding artifacts by reducing merging in dynamic regions and in tiles which may have been aligned incorrectly.

Our contributions on top of [Hasinoff et al. 2016] are 1) Calculating mismatch maps, in which high mismatch represents dynamic regions in the scene; 2) Using these maps to enable spatially varying Fourier-domain burst merging, in which merging is increased in static regions and reduced in dynamic regions; and 3) Increasing spatial denoising where merging was limited, to compensate for the reduced temporal denoising.

3.1 Alignment

Alignment and merging are performed with respect to a reference frame, chosen as the sharpest frame in the burst. We use a tile-based alignment algorithm similar to [Hasinoff et al. 2016] with slight modifications to improve performance in low light and high noise. The alignment is performed by a coarse-to-fine algorithm on four-level Gaussian pyramids of the raw input. In our implementation, the size of the tiles for alignment is increased as the noise level in the image rises, at the cost of slightly increased computation for larger tile sizes. Whereas the previous alignment method used 8 or 16 pixel square tiles, we use square tiles of 16, 32, and 64 pixels, depending on the noise level. For example, in a well-lit indoor setting, the size of the alignment tiles is 16 while in the dark it is 64 pixels.

3.2 Spatially varying temporal merging

As described in [Hasinoff et al. 2016], after alignment, tiles of size 16×16 pixels are merged by performing a weighted average in the Fourier domain. This merging algorithm is intrinsically resilient to temporal changes in the scene and to alignment errors, owing to the reduced weight given to frames and frequency bands that differ from the reference frame significantly above the noise variance. In [Hasinoff et al. 2016], the weight of the contribution from each frame z in the location of tile t , is proportional to $1 - A_{tz}(\omega)$, in which $A_{tz}(\omega)$ is

$$A_{tz}(\omega) = \frac{|D_t z(\omega)|^2}{|D_t z(\omega)|^2 + c\sigma_{tz}^2}, \quad (7)$$

where $D_t z(\omega)$ is the difference between tiles in Fourier band ω as defined in [Hasinoff et al. 2016], σ_{tz}^2 is the noise variance provided by a model of the noise, and c is a constant that accounts for the scaling of the noise variance. If the difference between tiles is small relative to the noise, $A_{tz}(\omega)$ decreases and a lot of merging occurs, and if the difference is large, the function goes to 1 and the contribution of the tile decreases.

In Equation 7, the value of c can be tuned to control the tradeoff between denoising strength and robustness to motion and alignment-error artifacts. We call this parameter the “temporal strength”, where a strength of 0 represents no temporal denoising, and a strength of infinity represents averaging without any motion robustness (see Figure 9a–d).

For very low-light scenes, we have found that c has to be increased considerably in order to increase the contribution of each frame to the average and to reduce noise. However, doing this introduces artifacts due to tiles erroneously being merged together (due to scene motion or alignment errors, as shown in Figure 9). In order to reduce noise in very low-light scenes without suffering from artifacts, we introduce a spatially varying scaling factor to c , the “temporal strength factor”, f_{tz} , where t and z represent the tile and

frame respectively. The modified temporal strength is $c \cdot f_{tz}$. We compute the temporal strength factor as a function of a mismatch map, m_{tz} , which is calculated per frame in the spatial (pixel) domain. The mismatch maps are calculated using a “shrinkage” operator [Dabov et al. 2007], with the same form as Equation 7:

$$m_{tz} = \frac{d_{tz}^2}{d_{tz}^2 + s\sigma_{tz}^2}. \quad (8)$$

The mismatch, m_{tz} , is calculated between the corresponding tile in frame z to tile t in the reference frame, and has a value between zero (no mismatch) and one (high mismatch). d_{tz} is the L1 difference between tile t in the reference frame and the tile that was aligned to it from frame z . For efficiency, we reuse the calculations in the previous alignment stage to obtain d_{tz} . The constant s is a scaling factor to allow tuning the mismatch maps, and we found that $s = 0.5$ results in a mismatch map that suppresses noise while detecting actual motion in the scene.

The temporal strength factor, f_{tz} , is a piecewise-linear function of m_{tz} , and the modified $A_{tz}(\omega)$ is now:

$$A_{tz}(\omega) = \frac{|D_{tz}(\omega)|^2}{|D_{tz}(\omega)|^2 + c f_{tz}(m_{tz}) \sigma_{tz}^2}. \quad (9)$$

The function $f_{tz}(m_{tz})$, shown in Figure 9f, is manually tuned for various scene noise levels (and linearly interpolated between them) to increase the temporal denoising in high-matching regions while decreasing merging and preserving the details in regions that have high-mismatch. The maximal value of the temporal strength factor is higher when more merging is desired, notably, darker scenes with higher noise, as determined by a calculated tuning factor. The minimal value of the temporal strength factor is similar for all scenes and represents increased robustness to mismatch artifacts. In these cases the temporal strength factor is simply 1 and the algorithm naturally degrades back to Equation 7.

This modified algorithm provides additional robustness to [Hasinoff et al. 2016] by better determining when the differences between tiles originate from changes in random noise and when they are due to actual changes in the scene. Although the mismatch map is also based on tile differences, the difference of Equation 7, $D_{tz}(\omega)$, only looks at differences within a certain frequency band in a 16×16 pixel tile, while the difference d_{tz} of Equation 8 originates from the alignment error which often has a larger support (up to 64×64) and is calculated in the spatial domain, across all frequency bands. These complementary approaches to calculating the differences between tiles help create a motion robust merge that is also effective at removing noise in very low light.

As an additional improvement to [Hasinoff et al. 2016], instead of calculating the weight $1 - A_{tz}(\omega)$ per color channel and merging each channel separately, we now compute the minimal weight across all color channels for each tile, and use a similar weight for merging all of the channels in each frequency band. We found that in low light the weights for each channel can be very different from one another, which causes color artifacts. Using a similar weight is important to reduce these artifacts.

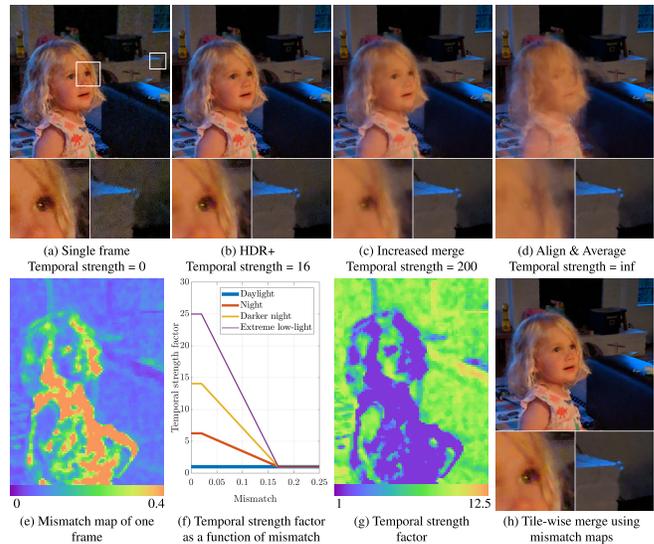


Fig. 9. The temporal strength (c in Equation 7) determines the tradeoff between temporal denoising and motion robustness. It does this by adjusting the contributions of tiles in alternate frames to the merged result based on their similarity to the corresponding tiles in the reference frame. To demonstrate the influence of the temporal strength, we manually set it to various values in sub-figures (a-d). In (a) the temporal strength is set to zero and effectively only the reference frame is used to create the image. In (b), the burst of 13 frames is merged using a low temporal strength. The resulting image has significant noise but no motion blur. In (c, d) the temporal strength is high and consequently the resulting image has less noise but more motion blur. By creating mismatch maps for each frame, we can compute a spatially varying temporal strength. (e) shows a mismatch map of one of the frames of the burst, as computed per tile according to Equation 8. The static background has low mismatch and the regions which have motion have a high mismatch. (f) shows the conversion from the mismatch maps to the temporal strength factor at various scene noise levels. (g) shows the spatially varying temporal strength factor. In the resulting image, regions with movement, such as the face, are sharp as in (b), and the static regions, such as the corner of the fireplace, have reduced noise as in (c). Refer to the supplement to see additional frames from this burst and the corresponding mismatch maps.

3.3 Spatially varying denoising

As in [Hasinoff et al. 2016], we next perform spatial denoising in the 2D DFT domain, by applying a shrinkage operator of the same form as Equation 7 to the spatial frequency coefficients. Originally, the strength of denoising was calculated assuming that all N frames were averaged perfectly and the estimate of the noise variance of the image after merging was updated to be $\frac{\sigma^2}{N}$. In our modified implementation, we calculate the noise variance of the merged image based on the actual merging that occurred in each tile, which increases the noise variance in regions of high mismatch. As a result, these regions get increased spatial denoising, via the shrinkage function, which compensates for their reduced temporal denoising. See Figure 10 for the results in Figure 9 without and with spatially varying denoising.



Fig. 10. Spatially varying denoising adaptively reduces more noise in regions with less temporal denoising, such as the face and hand. Static regions of the image are not affected. The amount of denoising increases in tiles with reduced frame merging.

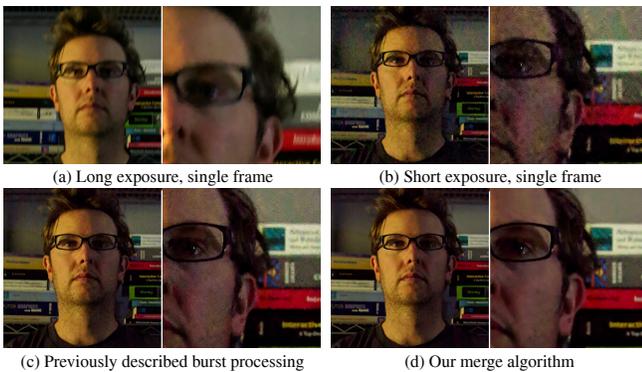


Fig. 11. Motion-robust align and merge can be applied to dSLR photos to improve low light handheld photography. a) At 0.3 lux, at an exposure time of 330 ms at ISO 12800, the photograph isn't overly noisy on a professional dSLR (Canon EOS-1D X) but suffers from motion blur. In these situations, photographers must balance a tradeoff between unwanted motion blur and unwanted noise. b) At 50 ms exposure and ISO 65535 there is no motion blur, but the image suffers from significant noise. c) Merging 13 frames with 50 ms exposure at ISO 65535 as described in [Hasinoff et al. 2016] produces an image with reduced noise and reduced motion blur. d) When applying our modified merge algorithm to the burst we are able to reduce noise even further and provide more detail. The crops on the right were brightened by 30% to emphasize the differences in noise.

Figure 9h shows the result of our algorithm on a dark indoor scene with a lot of motion. The image was created by merging a burst of 13 frames captured on a mobile device. Our procedure for aligning and robustly merging multiple frames generalizes to the raw bursts captured by a handheld digital single-lens reflex camera (dSLR). As shown in Figure 11, our algorithm can be used to merge multiple short exposures and reduce motion blur.

4 LOW-LIGHT AUTO WHITE BALANCE

Camera pipelines require an Automatic White Balance (AWB) step, wherein the color of the dominant illumination of the scene is estimated and used to produce an aesthetically pleasing color composition that (often) appears to be lit by a neutral illuminant. White

balance is closely related to color constancy, the problem in both human vision [Foster 2011] and computer vision [Gijssen et al. 2011] of discounting the illuminant color of a scene and inferring the underlying material colors. Because modern color constancy algorithms tend to use machine learning, these algorithms can be adapted to serve as AWB algorithms by training them on illuminants that, when removed, result in pleasing color-corrected images, rather than training them on the illuminants used in color constancy that, when removed, produce veridically neutralized images. “Fast Fourier Color Constancy” (FFCC) [Barron and Tsai 2017] is the current top-performing color constancy algorithm, both in terms of accuracy (on standard color constancy datasets [Cheng et al. 2014; Gehler et al. 2008]) and speed (capable of running at ~ 700 frames per second on a mobile device). FFCC works by learning and then applying a filter to a histogram of log-chrominance pixel and edge intensities, as was done by its predecessor technique [Barron 2015].

White balance is a particularly challenging problem in low-light environments, as the types of illuminations are more diverse, and the observed image is more noisy. In this section we describe how FFCC can be modified to serve as an effective AWB algorithm for low-light photography, through the use of a new dataset and a new error metric for training and evaluation. In the supplement we document some additional modifications we made to FFCC that are unrelated to low light photography, regarding image metadata and sensor calibration.

Dataset: FFCC is learning-based, and so its performance depends critically on the nature of the data used during training. While there are publicly available color constancy datasets [Cheng et al. 2014; Gehler et al. 2008], these target color constancy—objectively neutralizing the color of the illuminant of a scene. Such neutralized images often appear unnatural or unattractive in a photographic context, as the photographer often wants to retain some of the color of the illuminant in particular settings, such as sunsets or night clubs. Existing datasets also lack imagery in extreme low-light scenarios, and often lack spectral calibration information. To address these issues, we collected 5000 images of various scenes with a wide range of light levels, all using mobile devices with calibrated sensors. Rather than using a color checker or gray card to recover the “true” (physically correct) illuminant of each scene, we recruited professional photographers to manually “tag” the most aesthetically preferable white balance for each scene, according to their judgment. For this we developed an interactive labelling tool to allow a user to manually select a white point for an image, while seeing the true final rendering produced by the complete raw processing pipeline of our camera. Matching the on-device rendering during tagging is critical, as aesthetic color preferences can be influenced by other steps during image processing, such as local tone-mapping, exposure, and saturation. Our tagging tool and a subset of our training data is shown in the supplement.

Error metrics: Traditional error metrics used for color constancy behave erratically in some low-light environments, which necessitates that we design our own error metric for this task. Bright environments are often lit primarily by the sun, so the true illuminant of images taken in bright environments is usually close to white. In contrast, low light environments are often lit by heavily tinted illuminants—campfires, night clubs, sodium-vapor lamps, etc. This

appears to be due to a combination of sociological and technological trends (the popularity of colorful neon lights in dark environments) and of the nature of black body radiation (dim lights are often tinted red). Heavily-tinted illuminants mean that the true white-balanced image may contain pixel values where a color channel's intensity is near zero for all pixels, such as the image in Fig. 12a. This is problematic when evaluating or training white balance algorithms, which are commonly evaluated in terms of how well they recover *all* color channels of the illuminant. If a color channel is “missing” (i.e. having little or no intensity) then this task becomes ill-posed. For example, if all pixels in an image have a value near 0 in their red channel, then the image will look almost identical under all possible scalings of its red channel. This is problematic for a number of reasons: 1) When collecting ground-truth illuminant data, it is unclear how missing channels of the illuminant should be set. 2) When evaluating an algorithm on heavily-tinted images, its accuracy may seem low due to irrelevant “mistakes” when estimating missing color channels. 3) Learning-based algorithms will attempt to minimize a loss function with respect to meaningless illuminant colors in these missing channels, which may impair learning—a model may waste its capacity trying to accurately predict illuminant colors that are impossible to determine and are irrelevant to the image's final rendering.

To address this, we present a new error metric that correctly handles heavily-tinted scenes. First, let us review the canonical metric for color constancy, which is the angular error between RGB vectors of the true illuminant ℓ_t and the recovered illuminant ℓ_p [Gijssen et al. 2009]:

$$\Delta_\ell(\ell_p, \ell_t) = \cos^{-1} \left(\frac{\ell_p^T \ell_t}{\|\ell_p\| \|\ell_t\|} \right) \quad (10)$$

The angular error measures errors in the recovered illuminant, even though our primary concern is the appearance of the white-balanced image attained when dividing the input image by the recovered illuminant. The reproduction error of [Finlayson and Zakizadeh 2014] addresses this:

$$\Delta_r(\ell_p, \ell_t) = \cos^{-1} \left(\frac{\|r\|_1}{\sqrt{3}\|r\|_2} \right) \quad r = \frac{\ell_t}{\ell_p} \quad (11)$$

where division is element-wise. The reproduction error considers the error in the appearance of a white patch under a recovered illumination compared to the true illuminant. But as we have noticed, in heavily-tinted scenes where one or two color channels are missing, the notion of a “white patch” becomes ill-defined. We therefore propose an improved metric that considers the appearance of the *average* patch of the image when viewed under the recovered illumination, which we dub the “Anisotropic Reproduction Error” (ARE):

$$\Delta_a(\ell_p, \ell_t, \mu_t) = \cos^{-1} \left(\frac{\sqrt{r}^T H \sqrt{r}}{\sqrt{\text{tr}(H)} \sqrt{r^T H r}} \right) \quad H = \text{diag}(\mu_t)^2 \quad (12)$$

μ_t is an RGB vector containing the mean color of the *true* image:

$$\mu_t = \text{mean}_i \left(\begin{bmatrix} r_t^{(i)} \\ g_t^{(i)} \\ b_t^{(i)} \end{bmatrix} \right) \quad (13)$$

The ARE implicitly projects the inverse of the true and predicted illuminants (i.e. the appearance of a white patch under the true

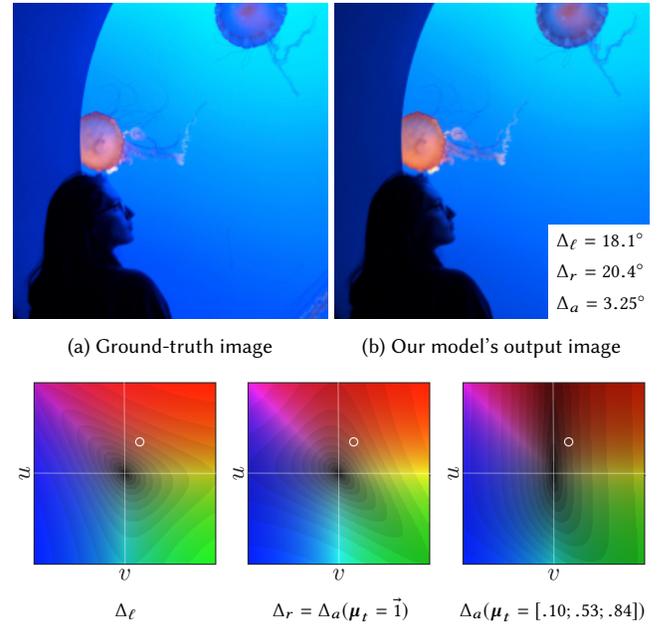


Fig. 12. Low-light environments often have heavily tinted illuminants, resulting in images where one or more color channels have low intensities across the image (a). As a result, during white balance the appearance of the image is sensitive to accurate estimation of the gains for high-intensity channels (here, green and blue) but is less sensitive to the low-intensity channels (here, red), as evidenced by the output of our model (b), which is grossly inaccurate in the red channel but still results in an accurate-looking rendering. Traditional error metrics (Δ_ℓ and Δ_r) exaggerate errors in our low-intensity color channel (b, inset), causing difficulties during training and evaluation. Our “Anisotropic Reproduction Error” (ARE), denoted Δ_a , accounts for the anisotropy of the RGB values in the true image, and down-weights errors in missing channels accordingly, resulting in a learned model that produces output images (b) that closely match the ground-truth (a) in these challenging scenes. Below we visualize the error surfaces of traditional metrics Δ_ℓ and Δ_r alongside Δ_a using the log-chroma UV histograms of [Barron and Tsai 2017], where luma indicates error (the true illuminant is at the origin and our prediction is a white circle). By conditioning on the mean RGB value of the true image μ_t , the ARE's curvature is reduced with respect to red (along the u -axis) but still highly curved with respect to blue and green (along the v -axis and diagonally), which accurately reflects performance in this heavily-tinted scene.

and predicted illuminants) into a 3D space scaled by the average color of the true white-balanced image μ_t , and then computes the reproduction error in that projected space. Therefore, the smaller the value of some channel of μ_t , the less sensitive Δ_a is to errors in that channel of the estimated illuminant ℓ_p , with a value of 0 in μ_t resulting in complete insensitivity to errors in that channel. If two color channels are absent in an image, the ARE will be 0 regardless of the predicted illuminant, accurately reflecting the fact that the color composition of the scene cannot be changed by applying per-channel color gains. The ARE degrades naturally to the standard reproduction error when the average scene color is gray:

$$\forall \alpha > 0, \Delta_a(\ell_p, \ell_t, \alpha \mathbf{1}) = \Delta_r(\ell_p, \ell_t) \quad (14)$$

Table 3. White balance results using 3-fold cross-validation on our dataset, in which we modify FFCC by replacing its data term with different error metrics, including the ARE, Δ_a . We report mean error, and the mean of the largest 25% of errors. Minimizing the ARE improves performance as measured by *all* error metrics, as it allows training to be invariant to the inherent ambiguity of white points in heavily-tinted scenes.

Algorithm	Δ_ℓ (Ang. Err.)		Δ_r (Repro. Err.)		Δ_a (ARE)	
	Mean	Worst 25%	Mean	Worst 25%	Mean	Worst 25%
FFCC [2017]	1.377	3.366	1.836	4.552	1.781	4.370
FFCC + Δ_r	1.332	3.154	1.771	4.255	1.721	4.088
FFCC + Δ_ℓ	1.320	3.115	1.768	4.229	1.714	4.049
FFCC + Δ_a	1.312	3.088	1.752	4.194	1.696	4.007

In most of the scenes in our dataset, the average true color is close to gray and therefore the error surface of the ARE closely resembles the reproduction error, but in challenging low-light scenes the difference between the two metrics can be significant, as is shown in Figure 12

In addition to its value as a metric for measuring performance, the ARE can be used as an effective loss during training. To demonstrate this we trained our FFCC model on our dataset using the same procedure as was used in [Barron and Tsai 2017]: three-fold cross validation, where hyperparameters are tuned to minimize the “average” error used by that work. We trained four models: a baseline in which we minimize the von Mises negative log-likelihood term used by [Barron and Tsai 2017], and three others in which we replaced the negative log-likelihood with three error metrics: Δ_r , Δ_ℓ , and our Δ_a (all of which are differentiable and therefore straightforward to minimize with gradient descent). Results can be seen in Table 3. As one would expect, minimizing the ARE during training produces a learned model that has lower AREs on the validation set. But perhaps surprisingly, we see that training with ARE produces a learned model that performs better on the validation data *regardless of the metric used for evaluation*. That is, even if one only cared about minimizing angular error or reproduction error, ARE is still a preferable metric to minimize during training. This is likely because the same ambiguity that ARE captures is also a factor during dataset annotation: If an image is lacking some color, the annotation of the ground-truth illuminant for that color will likely be incorrect. By minimizing ARE, we model this inherent uncertainty in our annotation and are therefore robust to it, and our learned model’s performance is not harmed by potentially-misleading ground truth annotations.

5 TONE MAPPING

The process of compressing the dynamic range of an image to that of the output is known as “tone mapping”. In related literature, this is typically achieved through the application of tone mapping operators (TMOs), which differ in their intent: some TMOs try to stay faithful to the human visual system, while others attempt to produce images that are subjectively preferred by artistic experts. [Eilertsen et al. 2016; Ledda et al. 2005] provide evaluations of TMOs for various intents.

Figure 13 depicts the human visual acuity in different light levels. The intent of our TMO is to render images that are vibrant and

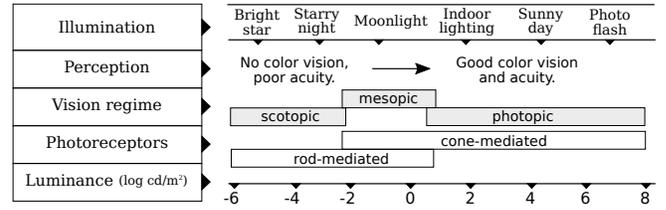


Fig. 13. The variation in color perception and acuity of human vision across typical ambient light levels (figure adapted from [Stockman and Sharpe 2006]). Any TMO that attempts to simulate the human visual system in low-light situations (scotopic and mesopic regimes) is subject to complex physiological phenomena, including Purkinje shift, desaturation, and loss of spatial acuity [Shin et al. 2004].

colorful even in low-light scenarios, enabling photopic-like vision at mesopic and scotopic light levels. While TMOs are able to simulate the human visual system in the color-sensitive cone mediated regime, simulating the rod-mediated regime is more challenging due to the presence of complex physiological phenomena [Ferwerda et al. 1996; Jacobs et al. 2015; Kirk and O’Brien 2011]. The ability to render vibrant and colorful dark scenes has been developed for digital astronomy [Villard and Levay 2002]. Our system aims to bring this ability to handheld mobile cameras.

Development of the final look of a bright nighttime photograph needs to be done carefully. A naïve brightening of the image can lead, for example, to undesired saturated regions and low contrast in the result. Artists have known for centuries how to evoke a nighttime aesthetic through the use of darker pigments, increased contrast, and suppressed shadows ([Helmholtz 1995], originally written in 1871). In our final rendition, we employ the following set of carefully tuned heuristics atop the tone mapping of [Hasinoff et al. 2016] to maintain a nighttime aesthetic:

- Allow higher overall gains to render brighter photos at the expense of revealing noise.
- Limit boosting of the shadows, keeping a significant amount of the darkest regions of the image near black to maintain global contrast and preserve a feeling of nighttime.
- Allow compression of higher dynamic ranges to better preserve highlights in challenging conditions, such as night scenes with light sources.
- Boost the color saturation in inverse proportion to scene brightness to produce vivid colors in low-light scenes.
- Add brightness-dependent vignetting.

[Hasinoff et al. 2016] uses a variant of synthetic exposure fusion [Mertens et al. 2007] that splits the full tonal range into a shorter exposure (“highlights”) and longer exposure (“shadows”), determines the gains that are applied to each exposure, and then fuses them together to compress the dynamic range into an 8-bit output. For scenes with illuminance E_v below log-threshold L_{\max} , we gradually increase the gain A_s applied to the shadows (the longer synthetic exposure) with the lux level, reaching a peak of over 1 additional stop (2.2×) of shadow brightness at log-threshold L_{\min} :

$$A_s = 2.2^{1 - \max\left(0, \min\left(1, \frac{\log(E_v) - L_{\min}}{L_{\max} - L_{\min}}\right)\right)}, \quad (15)$$



Fig. 14. A nighttime scene (about 0.4 lux) that demonstrates the difficulty in achieving balanced tonal quality. (a) The tone mapping of [Hasinoff et al. 2016] results in most of the image being too dark and is not a pleasing photograph. (b) Applying a tone mapping technique that brightens the image using histogram equalization (CLAHE [Zuiderveld 1994]) reveals more detail, but the photo lacks global contrast. (c) The photo with our tone mapping retains the details of (b), but contains more global contrast and enough dark areas to show that the captured scene is dark.

where the illuminance thresholds L_{\max} and L_{\min} correspond to a scene brightness of 200 lux and 0.1 lux, respectively.

A fraction of up to 20% of the additional shadow gain A_s is also applied to the highlights (the short synthetic exposure). This is done in inverse proportion to the normalized dynamic range of the scene D , because otherwise the intensity distribution can become too narrow, causing a loss of contrast. The highlight gain A_h is:

$$A_h = 1 + 0.2 \cdot (A_s - 1) \cdot (1 - D). \quad (16)$$

We also increase the color saturation in inverse proportion to the log-lux level of the scene up to 20%, which helps maintain vivid renditions of facial skin tones.

Due to the high gains applied to the images, artifacts from inaccuracies in the subtraction of the black-level and lens-shading correction begin to manifest. These artifacts are most prominent in the corners of the image (where the lens-shading correction gain is greatest). To suppress these artifacts, we apply vignetting in scenes darker than 5 lux, gradually increasing in darker scenes. The appearance of vignetting also contributes to the perception of a dark scene and adds an aesthetic “depth” to the image.

To demonstrate the results of our method, the rendition from [Hasinoff et al. 2016] is shown in Figure 14a and is too dark to show much detail. The amount of detail in that rendition can be increased by locally increasing its contrast, for example by applying adaptive histogram equalization. In Figure 14b we use CLAHE [Zuiderveld 1994] on the V channel following an HSV decomposition of the image [Smith 1978]. This increases local contrast but the global contrast is flattened, resulting in a low-contrast appearance. The result of our technique (Figure 14c) reveals shadow details without significant loss of contrast and maintains a natural nighttime appearance.

6 RESULTS

Our system was launched publicly in November 2018 as a mode called “Night Sight” in the camera app on Google Pixel smartphones. It is implemented on Android and relies on the Camera2 API [Google

LLC 2016a]. On Pixel, Pixel 2 and Pixel 3a we use the motion-adaptive burst merging method described in Section 3. On Pixel 3, which has a faster system-on-chip, burst merging is performed using [Wronski et al. 2019]. This alternative merging technology also launched in October 2018 on the Pixel 3 under the name “Super Res Zoom”. We compare these two merging algorithms in the supplementary material. All of the results in this paper, except for those clearly indicated in the supplementary comparison figure, use the motion-robust Fourier-domain merge described in Section 3.

Our system uses refined and tuned versions of some of the algorithms described in the “Finishing” section of [Hasinoff et al. 2016] in order to convert the raw result of the merged burst into the final output. These include demosaicing, luma and chroma spatial denoising, dehazing, global tone adjustment, and sharpening. We are not able to provide additional details about these proprietary algorithms, and we note that they are not a contribution of the system described in this paper.

In Figure 15, we compare our results to the burst photography pipeline that we build upon [Hasinoff et al. 2016]. The sequence shows that simply applying the tone mapping described in Section 5 to the baseline results in excessive noise. When motion metering (Section 2) and motion-adaptive burst merging (Section 3) are applied, the noise and detail improve significantly. Finally, the low-light auto white balance (Section 4) algorithm improves color rendition, making the final image better match the human visual perception of the scene. The final results allow users to see more detail while still having low noise, pleasing colors, and good global contrast. In the supplement, we demonstrate our processing compared to [Hasinoff et al. 2016] across multiple scenes and a variety of illuminants. Additionally, we have processed the dataset released by [Hasinoff et al. 2016] ([Google LLC 2016b]) and published our results on [Google LLC 2019].

We also compare our pipeline with an end-to-end trained neural network that operates on raw images [Chen et al. 2018]. The CNN performs all of the post-capture tasks performed by our pipeline, but with a large convolutional architecture and a single frame as input. Figure 16 shows that the output of the CNN has less detail and more artifacts than the output of our pipeline (additional results are in the supplement). In this comparison, we used raw frames captured with a similar camera model as was used for training the network in [Chen et al. 2018] (Sony $\alpha 7S$ II) because that network did not generalize well to the frames of the mobile device and produced severe color and detail artifacts (not shown here). Although our pipeline was not specifically calibrated or tuned for the Sony camera, it still produced higher-quality images compared with the results of [Chen et al. 2018].

Performance on mobile devices must also be considered while evaluating imaging pipelines. The CNN network architecture in [Chen et al. 2018] is large and may not be well suited for running efficiently on a mobile device. [Chen et al. 2018] reports a time of 0.38 s to process the Sony images on a desktop GPU, which is equivalent to 10s of seconds on a mobile GPU [GFXBench 2019] (assuming memory and temperature limits allow the model to be run on a mobile device). Our pipeline has been designed for efficiency, and thus is able to process the captured burst of images in under 2 seconds on a mobile phone. Pre-capture motion metering is also efficient and



Fig. 15. Comparison between our pipeline and the baseline burst processing pipeline [Hasinoff et al. 2016] on 13 frame bursts. Due to hand shake, over the duration of the capture the top example contains 41 pixels of motion and the bottom example contains 66 pixels of motion. The columns show the effect of augmenting the baseline with various parts of our pipeline. (a) shows the results of the baseline. (b) demonstrates how our tone mapping makes the scene brighter, allowing the viewer to see more detail while preserving global contrast. However, noise becomes more apparent, and overwhelms visible detail in the darker parts of the scenes. Adding motion metering in (c) allows us to increase the exposure time in these scenes without introducing significant motion blur (3-4 pixels per frame). The resulting images are less noisy, and more detail is visible. Adding motion-robust merging in (d) allows us to increase temporal denoising in static regions, further improving detail. Finally, the low-light auto white balance is applied in (e), making the colors in the scene better match the perceived colors from the human visual system.



Fig. 16. Comparison between the end-to-end trained CNN described in [Chen et al. 2018] and the pipeline in this paper. Raw images were captured with a Sony $\alpha 7S$ II, the same camera that was used to train the CNN. The camera was stabilized using a tripod. The illuminance of these scenes was 0.06-0.15 lux. Because our white balancing algorithm was not calibrated for this sensor, and since color variation can affect photographic detail perception, the results of our pipeline are color-matched to the results of [Chen et al. 2018] using Photoshop's automatic tool (see supplement for details). Column (a) shows the result from the CNN, and column (b) shows the result of our pipeline processing 13 raw frames. Column (c) shows the result of our pipeline processing one raw frame (no merging). The advantages of burst processing are evident, as the crops in column (b) show significantly more detail and less noise than the single-frame results. Additional examples can be found in the supplement.



Fig. 17. Comparison of the output of our system to photographs captured with the specialized low-light modes of other devices. The scene brightness is 0.4 lux and it is illuminated by a dim tungsten light. Readers are encouraged to zoom in on the images. The different results show the challenges inherent in low-light photography. High levels of detail, limited artifacts, pleasing color, and good exposure (brightness) are all important, and are difficult to achieve in unison.

can estimate the flow field’s magnitude in only 3 ms. Therefore, the maximum total capture and processing time of our system is 8 s, of which 6 s is the maximum capture time, achieved when the camera is stabilized. Memory is also an important consideration. Based on the architecture of the network described in [Chen et al. 2018], we estimate that it requires around 3 GB of memory, while our process needs under 400 MB of memory.

Figure 17 compares photos captured with our device and processed with our on-device pipeline, to three other devices that have specialized low-light photography modes. Our pipeline has a well-balanced rendition that has good brightness and pleasing color with a low amount of artifacts. For additional examples, please refer to the supplement.

6.1 Limitations

Although our proposed system is successful in producing high-quality photographs on a mobile device in low-light conditions, it has several practical limitations.

Autofocus. Consistent and accurate autofocus (AF) is a prerequisite for sharp output images. AF in low-SNR scenes is fundamentally difficult, and beyond the scope of this paper. We use the device’s AF out of the box, which employs both phase- and contrast-based strategies, but isn’t optimized for very low-light conditions. Empirical experiments show that our camera’s AF fails almost 50% of the time when the scene luminance is around 0.4 lux. We mitigate this issue by locking the lens’s focus distance to the hyperfocal distance when the AF system reports a focus failure. When the lens is at the hyperfocal distance, everything from a certain distance out to infinity is in acceptable focus, and the depth of field is maximized,

thereby minimizing potential misfocus. We also allow the user to adjust the focus distance manually prior to the capture.

Real-time preview. The viewfinder allows the user to frame the shot in real time and anticipate the final rendition of a processed image. Rendering a high-quality preview in low light is challenging, as both exposure time and processing latency are constrained by the viewfinder refresh rate (e.g., ≥ 15 fps). We circumvent this challenge by adopting the camera’s proprietary hardware-accelerated preview pipeline and adaptively gaining up its output using a GPU shader. One could potentially obtain better results by employing lightweight temporal denoising techniques such as [Ehmann et al. 2018].

Shutter lag. Similar to [Hasinoff et al. 2016], our system doesn’t start collecting raw frames until the user presses the shutter, resulting in a noticeable delay for the captured “moment”. This is different from many zero-shutter-lag (ZSL) applications, where the camera continuously caches viewfinder frames to a ring buffer and processes them upon shutter press. We choose not to use ZSL frames because their exposure time is limited by the viewfinder refresh rate, which is so short that read noise dominates the signal.

Black levels. The minimum digital value returned by a sensor that is treated as black is called the black level. Inaccurate black levels are a potential limitation to low-light imaging. However, a good sensor can get within 0.25 data numbers (DN) of a 10-bit signal, and accuracy can be further improved with a gain-dependent look-up table. We use black levels straight from the camera driver.

Sky tones. At night, the tone of the sky usually overlaps with that of the foreground. As our pipeline adjusts the mid-tones during global tone mapping to make details in the foreground more visible, the rendition of the sky is also brightened. The perception of darkness is naturally sensitive to the tone and color of the sky and therefore in certain conditions our results fail to convey the sense of darkness adequately. To correctly account for the perception of darkness, one would need to apply a different tonal adjustment specifically to the sky.

7 CONCLUSIONS

In this paper, we set forth a list of requirements that a computational photography system should satisfy in order to produce high-quality pictures in very low-light conditions, and described an implementation of a system that meets these requirements. We introduced motion metering, a novel approach for determining burst capture parameters that result in a good tradeoff between SNR and motion blur, and motion-robust merging, which minimizes motion artifacts when merging extremely noisy raw images. We also outlined how we incorporate learning-based auto white balancing and specialized tone mapping to best convey the color and aesthetics of a low-light photograph. Our system overcomes many constraints of a small sensor size, and runs efficiently on a modern mobile device.

ACKNOWLEDGMENTS

Night Sight and the algorithms presented in this paper would not have been possible without close collaboration with Google Research’s Gcam team, Luma team, and Android camera team. From Gcam we particularly thank David E. Jacobs, Zalman Stern, Alexander Schifffhauer, Florian Kainz, and Adrian Rapazzini for their help

on the project and paper, and staff photographers Michael Milne, Andrew Radin, and Navin Sarma. From the Luma team, we particularly thank Peyman Milanfar, Bartłomiej Wronski, and Ignacio Garcia-Dorado for integrating and tuning their super-resolution algorithm for Night Sight on Pixel 3 and for their helpful discussions.

REFERENCES

- Miika Aittala and Frédo Durand. 2018. Burst image deblurring using permutation invariant convolutional neural networks. *ECCV* (2018).
- Jonathan T. Barron. 2015. Convolutional Color Constancy. *ICCV* (2015).
- Jonathan T. Barron and Yun-Ta Tsai. 2017. Fast Fourier Color Constancy. *CVPR* (2017).
- Giacomo Boracchi and Alessandro Foi. 2012. Modeling the Performance of Image Restoration From Motion Blur. *IEEE TIP* (2012).
- Tim Brooks, Ben Mildenhall, Tianfan Xue, Jiawen Chen, Dillon Sharlet, and Jonathan T. Barron. 2019. Unprocessing Images for Learned Raw Denoising. *CVPR* (2019).
- Daniel J. Butler, Jonas Wulff, Garrett B. Stanley, and Michael J. Black. 2012. A naturalistic open source movie for optical flow evaluation. *ECCV* (2012).
- Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. 2018. Learning to See in the Dark. *CVPR* (2018).
- Dongliang Cheng, Dilip K. Prasad, and Michael S. Brown. 2014. Illuminant estimation for color constancy: why spatial-domain methods work and the role of the color distribution. *JOSA A* (2014).
- Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. 2007. Image denoising by sparse 3-D transform-domain collaborative filtering. *TIP* (2007).
- Mauricio Delbracio and Guillermo Sapiro. 2015. Hand-Held Video Deblurring Via Efficient Fourier Aggregation. *TCI* (2015).
- Arthur P. Dempster, Nan M. Laird, and Donald B. Rubin. 1977. Maximum Likelihood from Incomplete Data Via the EM Algorithm. *Journal of the Royal Statistical Society, Series B*.
- Alexey Dosovitskiy, Philipp Fischery, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick van der Smagt, Daniel Cremers, and Thomas Brox. 2015. FlowNet: Learning Optical Flow with Convolutional Networks. *ICCV* (2015).
- Jana Ehmman, Lun-Cheng Chu, Sung-Fang Tsai, and Chia-Kai Liang. 2018. Real-Time Video Denoising on Mobile Phones. *ICIP* (2018).
- Gabriel Eilertsen, Jonas Unger, and Rafal K Mantiuk. 2016. Evaluation of tone mapping operators for HDR video. *High Dynamic Range Video* (2016).
- Sina Farsiu, M. Dirk Robinson, Michael Elad, and Peyman Milanfar. 2004. Fast and robust multiframe super resolution. *IEEE TIP* (2004).
- James A. Ferwerda, Sumanta N. Pattanaik, Peter Shirley, and Donald P. Greenberg. 1996. A model of visual adaptation for realistic image synthesis. *Computer graphics and interactive techniques* (1996).
- Graham Finlayson and Roshanak Zakizadeh. 2014. Reproduction Angular Error: An Improved Performance Metric for Illuminant Estimation. *BMVC* (2014).
- David H. Foster. 2011. Color constancy. *Vision research* (2011).
- Peter Vincent Gehler, Carsten Rother, Andrew Blake, Tom Minka, and Toby Sharp. 2008. Bayesian color constancy revisited. *CVPR* (2008).
- GFXBench. 2019. <http://gfxbench.com/>. [Online; accessed 17-May-2019].
- Arjan Gijsenij, Theo Gevers, and Marcel P. Lucassen. 2009. Perceptual analysis of distance measures for color constancy algorithms. *JOSA A* (2009).
- Arjan Gijsenij, Theo Gevers, and Joost van de Weijer. 2011. Computational Color Constancy: Survey and Experiments. *IEEE TIP* (2011).
- Clément Godard, Kevin Matzen, and Matt Uyttendaele. 2018. Deep Burst Denoising. *ECCV* (2018).
- Google LLC. 2016a. Android Camera2 API, <http://developer.android.com/reference/android/hardware/camera2/package-summary.html>.
- Google LLC. 2016b. HDR+ burst photography dataset, <http://www.hdrplusdata.org>.
- Google LLC. 2019. Handheld Mobile Photography in Very Low Light webpage, <https://google.github.io/night-sight/>.
- Yulia Gryaditskaya, Tania Pouli, Erik Reinhard, Karol Myszkowski, and Hans-Peter Seidel. 2015. Motion Aware Exposure Bracketing for HDR Video. *Computer Graphics Forum (Proc. EGSR)* (2015).
- Samuel W Hasinoff, Dillon Sharlet, Ryan Geiss, Andrew Adams, Jonathan T. Barron, Florian Kainz, Jiawen Chen, and Marc Levoy. 2016. Burst photography for high dynamic range and low-light imaging on mobile cameras. *SIGGRAPH Asia* (2016).
- Hermann von Helmholtz. 1995. On the relation of optics to painting. *Science and culture: Popular and philosophical essays* (1995).
- Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. 2017. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* (2017).
- Michal Irani and Shmuel Peleg. 1991. Improving resolution by image registration. *CVGIP: Graphical Models and Image Processing* (1991).
- David E. Jacobs, Orazio Gallo, Emily A. Cooper, Kari Pulli, and Marc Levoy. 2015. Simulating the Visual Experience of Very Bright and Very Dark Scenes. *ACM Trans. Graph.* (2015).
- Alexandre Karpenko, David Jacobs, Jongmin Baek, and Marc Levoy. 2011. Digital video stabilization and rolling shutter correction using gyroscopes. *Stanford Tech Report* (2011).
- Adam G. Kirk and James F. O'Brien. 2011. Perceptually based tone mapping for low-light conditions. *ACM Trans. Graph.* (2011).
- Patrick Ledda, Alan Chalmers, Tom Troscianko, and Helge Seetzen. 2005. Evaluation of tone mapping operators using a high dynamic range display. *SIGGRAPH* (2005).
- Marc Levoy. 2012. SynthCam, <https://sites.google.com/site/marclevoy/>.
- Ce Liu. 2009. *Beyond pixels: exploring new representations and applications for motion analysis*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- Ziwei Liu, Lu Yuan, Xiaoou Tang, Matt Uyttendaele, and Jian Sun. 2014. Fast Burst Images Denoising. *SIGGRAPH Asia* (2014).
- Bruce D. Lucas and Takeo Kanade. 1981. An Iterative Image Registration Technique with an Application to Stereo Vision. *IJCAI* (1981).
- Lindsay MacDonald. 2006. *Digital heritage*. Routledge.
- Matteo Maggioni, Giacomo Boracchi, Alessandro Foi, and Karen Egiazarian. 2012. Video denoising, deblocking, and enhancement through separable 4-D nonlocal spatiotemporal transforms. *TIP* (2012).
- Tom Mertens, Jan Kautz, and Frank Van Reeth. 2007. Exposure Fusion. *Pacific Graphics* (2007).
- Ben Mildenhall, Jonathan T. Barron, Jiawen Chen, Dillon Sharlet, Ren Ng, and Robert Carroll. 2018. Burst Denoising with Kernel Prediction Networks. *CVPR* (2018).
- Junichi Nakamura. 2016. *Image sensors and signal processing for digital still cameras*. CRC press.
- Shahriar Negahdaripour and C-H Yu. 1993. A generalized brightness change model for computing optical flow. *ICCV* (1993).
- Travis Portz, Li Zhang, and Hongrui Jiang. 2011. High-quality video denoising for motion-based exposure control. *IEEE International Workshop on Mobile Vision* (2011).
- Dr Rawat and Singhai Jyoti. 2011. Review of Motion Estimation and Video Stabilization techniques For hand held mobile video. *Signal & Image Processing: An International Journal (SIPIJ)* (2011).
- Jerome Revaud, Philippe Weinzaepfel, Zaid Harchaoui, and Cordelia Schmid. 2015. EpicFlow: Edge-Preserving Interpolation of Correspondences for Optical Flow. *CVPR* (2015).
- Erik Ringaby and Per-Erik Forssén. 2014. A virtual tripod for hand-held video stacking on smartphones. *ICCP* (2014).
- Simon Schulz, Marcus Grimm, and Rolf-Rainer Grigat. 2007. Using brightness histogram to perform optimum auto exposure. *WSEAS Transactions on Systems and Control* (2007).
- Jae Chul Shin, Hirohisa Yaguchi, and Satoshi Shioiri. 2004. Change of color appearance in photopic, mesopic and scotopic vision. *Optical review* (2004).
- Alvy Ray Smith. 1978. Color Gamut Transform Pairs. *SIGGRAPH* (1978).
- Andrew Stockman and Lindsay T Sharpe. 2006. Into the twilight zone: the complexities of mesopic vision and luminous efficiency. *Ophthalmic and Physiological Optics* (2006).
- Ray Villard and Zolta Levay. 2002. Creating Hubble's Technicolor Universe. *Sky and Telescope* (2002).
- Bartłomiej Wronski, Ignacio Garcia-Dorado, Manfred Ernst, Damien Kelly, Michael Krainin, Chia-Kai Liang, Marc Levoy, and Peyman Milanfar. 2019. Handheld Multi-Frame Super-Resolution. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2019).
- Karel Zuiderveld. 1994. Contrast limited adaptive histogram equalization. In *Graphics gems IV*.