



High-throughput fluorescence microscopy using multi-frame motion deblurring

ZACHARY F. PHILLIPS,^{1,3}  SARAH DEAN,^{2,3} BENJAMIN RECHT,²
AND LAURA WALLER^{1,2,*} 

¹Graduate Group in Applied Science and Technology, University of California, Berkeley, Berkeley, CA 94720-1700, USA

²Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, Berkeley, CA 94720-1700, USA

³These authors contributed equally to this work

*waller@berkeley.edu

Abstract: We demonstrate multi-frame motion deblurring for gigapixel wide-field fluorescence microscopy using fast slide scanning with coded illumination. Our method illuminates the sample with multiple pulses within each exposure, in order to introduce structured motion blur. By deconvolving this known motion sequence from the set of acquired measurements, we recover the object with up to 10× higher SNR than when illuminated with a single pulse (strobed illumination), while performing acquisition at 5× higher frame-rate than a comparable stop-and-stare method. Our coded illumination sequence is optimized to maximize the reconstruction SNR. We also derive a framework for determining when coded illumination is SNR-optimal in terms of system parameters such as source illuminance, noise, and motion stage specifications. This helps system designers to choose the ideal technique for high-throughput microscopy of very large samples.

© 2019 Optical Society of America under the terms of the [OSA Open Access Publishing Agreement](#)

1. Introduction

High-throughput wide-field microscopy enables the collection of large amounts of image data at high-speed, using optimized hardware and computational techniques to push system throughput beyond conventional limits. These systems play a critical role in drug discovery [1–3], functional protein analysis [4,5] and neuropathology [6–8], enabling the rapid acquisition of large datasets. In wide-field microscopes, the choice of objective lens defines both the resolution and the field-of-view (FOV) of the system, requiring the user to allocate optical throughput to *either* high-resolution *or* a wide FOV. Starting with a high-resolution objective, the FOV may be enlarged by mechanical scanning and image stitching, as in commercial slide-scanning systems [9]. Alternatively, computational imaging techniques have used a large FOV objective and enhanced the resolution of the system beyond the objective's wide-field diffraction limit [10–15].

Despite their wide adoption for a large variety of high-content imaging tasks, the performance of slide-scanning systems is often limited by the mechanical parameters of the motion stage rather than the optical parameters of the microscope [16]. The information throughput of an imaging system can be quantified by the space-bandwidth product (SBP), which is the dimensionless product of the spatial coverage (FOV) and Fourier coverage (resolution) of a system [17], as well as the space-bandwidth rate (SBR), which is SBP per unit time. Improving the SBP and SBR has been the subject of seminal works in structured illumination [13], localization microscopy [11,12], both conventional [14] and Fourier [10,15,18] ptychography, and deep learning [19–22]. Additionally, line-scan [23] and time-delay integration [24] methods use 1D sensors to reduce read-out time and increase SNR, but require specialized imaging hardware and calibration.

Quantifying the SBR of high-throughput imaging systems reveals bottlenecks in their acquisition strategy. For example, conventional wide-field slide-scanners are often SBR-limited by the time required for a motion stage to mechanically stabilize between movements. These mechanical

motions can lead to long acquisition times, especially when imaging very large samples such as coronal sections of the human brain at cellular resolution [25]. Conversely, a super-resolution technique such as Fourier ptychography only requires electronic scanning of LED illumination, so is more likely to be SBR-limited by photon counts or camera readout, since the time to change LED patterns is on the order of microseconds. Practically, the maximum resolution improvement is limited by the light-throughput at high illuminations [26] and the FOV is set by the optics.

Conventional slide-scanning microscopes employ one of two imaging strategies. The first, referred to as "stop-and-stare," involves moving the sample to each scan position serially, halting the stage motion before each exposure and resuming motion only after the exposure has finished. While this method produces high-quality images since it uses long exposures, it is slow due to the time required to stop and start motion between exposures. A second approach, referred to as "strobed", involves illuminating a sample in continuous motion with very short, bright pulses, in order to avoid the motion blur which would otherwise be introduced by an extended pulse. Strobed illumination will generally produce images with much lower SNR than stop-and-stare due to the short pulse times (often on the order of micro-seconds). The choice between these two acquisition strategies thus requires the user to trade-off SNR for acquisition rate, often in ways which make large-scale imagery impractical due to extremely long acquisition times. It should be noted that line-scan imaging systems are also subject to these trade-offs.

In this work, we propose a computational imaging technique which employs a coded-illumination acquisition for high-throughput applications. Our method involves continuously moving the sample (thus maintaining a high acquisition rate), while illuminating with multiple pulses during each acquisition (in order to achieve good SNR). This motion-multiplexing technique enhances the measurement SNR of our system, as compared to strobed illumination, by increasing the total amount of illumination. The resulting captured images contain motion-blur artifacts, which can be removed computationally through a multi-frame motion deblurring algorithm that uses knowledge of the pulse sequence and motion trajectory. The overall gain in SNR is proportional to the number of pulses as well as the conditioning of the motion deblurring process, necessitating careful design of pulse sequences to produce the highest-quality image.

In the following sections, we detail the joint design of the hardware and algorithms to enable gigapixel-scale fluorescence imaging with improved SNR (Fig. 1), compare the performance of our




Fig. 1. High-throughput microscope with temporally-coded illumination. A) Our system consists of an inverted wide-field fluorescence microscope with a 2-axis motion stage and a programmable LED illumination source. The sample is scanned at a constant speed while being illuminated with temporally coded illumination pulses during each exposure. The excitation wavelength filter may be removed for conventional brightfield imaging. B) Illustration of the pipeline from sample and acquired images to multi-frame reconstruction. C) Image of our system - a Nikon TE300 microscope configured with a Prior motion stage and LED illuminator [27].

proposed framework against traditional methods, and provide an experimental demonstration of situations where coded illumination is both optimal (e.g. fluorescence imaging) and sub-optimal (e.g. brightfield imaging) as a function of common system parameters such as illumination power and camera noise levels. Our contribution is both the proposal of a new high-throughput imaging technique as well as an analysis of when it is practically useful for relevant applications.

2. Methods

2.1. Motion blur forward models

With knowledge of both sample trajectory and illumination sequence, the physical process of capturing a single-frame measurement of a sample illuminated by a coded sequence of pulses while in motion can be mathematically described as a convolution:

$$\mathbf{y} = \mathbf{h} * \mathbf{x} + \eta, \tag{1}$$

where \mathbf{y} is the blurred measurement, \mathbf{x} is the static object to be recovered, η is additive noise, $*$ denotes 2D convolution, and \mathbf{h} is the blur kernel, which maps the temporal illumination intensity pattern to positions in the imaging coordinate system using kinematic motion equations.

For large FOV imaging, the sample is in continuous motion and the camera captures multiple frames. Hence, we must extend the single-frame coded-illumination forward model above to the multi-frame case. Mathematically, we model the motion-blurred frames as the vertical concatenation of many single-frame forward models, each of which are convolutional (Fig. 2). Each captured image has an associated blur operator \mathbf{B}_j defined by each blur kernel \mathbf{h}_j such that $\mathbf{h}_j * \mathbf{x} = \mathbf{B}_j \mathbf{x}$. Additionally, we prepend each convolutional sub-unit with a crop operator \mathbf{W}_j , which selects an area of the object based on the camera FOV. Together, these operators encode both spatial coverage and the local blurring of each measurement, and are concatenated to form the complete multi-frame forward operator, which is related to the measurements by:

$$\begin{bmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_n \end{bmatrix} = \begin{bmatrix} \mathbf{W}_1 \mathbf{B}_1 \\ \vdots \\ \mathbf{W}_n \mathbf{B}_n \end{bmatrix} \mathbf{x} + \eta = \mathbf{A} \mathbf{x} + \eta. \tag{2}$$




Fig. 2. Multi-frame motion deblurring forward model. \mathbf{y} is the blurred measurement and \mathbf{x} is the static object to be recovered. The operations in the blue box can be represented as a 1D convolution matrix \mathbf{A} , consisting of windowing operators \mathbf{W} and blur kernels \mathbf{h} . $[\cdot]$ denotes 2D convolution and $[\cdot]$ denotes the element-wise product.

This forward operator \mathbf{A} is no longer a simple convolution, but rather, a spatially-variant convolution based on the coverage of each individual crop operator. A 1D illustration of the multi-frame smear matrix \mathbf{A} is shown in Fig. 2. This forward operation and its adjoint can be computed efficiently within each crop window using the Fourier Transform.

2.2. Reconstruction algorithm

To invert our forward model (Eq. (2)), we employ the Nesterov accelerated gradient descent [28] algorithm to minimize the difference between our measurements \mathbf{y} and estimated object $\hat{\mathbf{x}}$ passed through forward model \mathbf{A} in the ℓ_2 metric. Here, we seek to minimize an unregularized cost function, $O(\mathbf{x}) = \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2$ using the following update equations at each k^{th} iteration:

$$\begin{aligned} \mathbf{z}^{k+1} &= \mathbf{x}^k - \alpha \mathbf{A}^H (\mathbf{A}\mathbf{x}^k - \mathbf{y}) \\ \mathbf{x}^{k+1} &= \beta^k \mathbf{z}^k + (1 - \beta^k) \mathbf{z}^{k+1}. \end{aligned} \quad (3)$$

α is a fixed step size, β^k is set each iteration by the Nesterov update equation, and \mathbf{z} is an intermediate variable to simplify the expression.

All reconstruction use 30 iterations of Eq. (3), which gives a favorable balance between reconstruction quality and reconstruction time. While adding a regularization term to enforce signal priors could improve reconstruction quality (and was previously analyzed in the context of motion deblurring [29]), we chose not to incorporate regularization in order to provide a more fair and straightforward comparison between the proposed coded-illumination and conventional strobed and stop-and-stare acquisitions. We instead investigate the role of regularization separately in Fig. 6. It should be noted that running our algorithm for a pre-defined number of iterations may provide some regularization from early-stopping [30].

Reconstructions were performed in Python using the Arrayfire GPU computation library [31]. Due to the raster-scanning structure of our motion pathway (Fig. 1(a)), the algorithm is highly parallelizable – hence, we separate our reconstruction into strips along the major translation axis and stitch these strips together after computation. Using this parallelization, we are able to reconstruct approximately 1 gigapixel in 2 minutes (further details in Section 3).

2.3. Reconstruction SNR

The primary benefit of coded illumination over strobed illumination is an improvement in signal strength. However, the deconvolution process amplifies noise, which can introduce artifacts in the reconstruction that negate these benefits. Hence, it is important to analyze the SNR of our system carefully for different system parameters.

To minimize the error between the reconstructed $\hat{\mathbf{x}}$ and the true object \mathbf{x} , the blur kernels and scanning pattern should be chosen such that the noise is minimally amplified by the inversion process. This amplification is controlled by the singular values of the forward model \mathbf{A} . In the case of single-frame blur, the singular values depend on the length and code of the blur kernel \mathbf{h} . While early works used a non-linear optimization routine (i.e. the `fmincon` function in MATLAB (Mathworks)) to minimize the condition number of \mathbf{A} [32,33], more recent work proposed maximizing the reconstruction SNR directly, using camera noise parameters, source brightness, and the well-posedness of the deconvolution [34,35]. We extend these works to the multiframe setting. In our analysis, we define SNR as the ratio of the mean signal to the signal variance (due to photon shot noise, camera readout noise, fixed pattern noise, and other camera-dependent factors). Under a simplified model, the noise variance will be the addition in quadrature of the camera read noise variance σ_r^2 plus a signal-dependent term \bar{s} :

$$SNR = \frac{\bar{s}}{\sqrt{\bar{s} + \sigma_r^2}}. \quad (4)$$

We ignore exposure-dependent noise parameters such as dark current and fixed-pattern noise, since these are usually small (for short exposure times) relative to read noise σ^2 . Note that the denominator of the SNR expression in Eq. (4) is equivalent to the standard deviation of additive noise, η , as in Eq. (1). This definition is valid for both strobed and stop-and-stare acquisitions.

For coded illumination acquisitions, it is necessary to consider the noise amplification that results from inverting the forward model. This amplification is controlled by the deconvolution noise factor (DNF) [34], which for single-frame blurring is defined as:

$$\text{DNF} = \sqrt{\frac{1}{m} \sum_{i=0}^m \frac{\max_i |\tilde{\mathbf{h}}|_i^2}{|\tilde{\mathbf{h}}|_i^2}}, \quad (5)$$

where m is the size of the blur kernel \mathbf{h} , and $\tilde{\mathbf{h}}$ represents the Fourier transform of \mathbf{h} .

We start by defining the single-frame SNR under coded illumination. To do so, we define a multiplexing factor $\gamma = \sum_i h_i$, the total amount of illumination imparted during exposure. If \mathbf{h} is constrained to be binary, γ will be equal to the total number of pulses. Equation (4) can then be modified using both the DNF and the multiplexing factor, γ :

$$\text{SNR} = \frac{\gamma \bar{s}_0}{\text{DNF} \sqrt{\gamma \bar{s}_0 + \sigma_r^2}}, \quad (6)$$

where \bar{s}_0 is the mean signal imparted by a single illumination pulse. This expression is valid for any additive noise model that is spatially uncorrelated. A full derivation of Eq. (6) is in Appendix 5.1, where we also discuss regularization and extensions to other noise models.

The derivation of the coded single-frame SNR in Eq. (6) relies on properties of convolutions, so the expression does not directly apply to the multi-frame forward operator \mathbf{A} , which is a spatially-variant convolution matrix. While methods for analyzing spatially-variant convolution matrices exist [36], the resulting expressions are complicated due to the boundary effects between frames. Therefore, we instead use a practical simplifying assumption: the blur path and illumination patterns are fixed to be the same across all frames, i.e. $\mathbf{h}_j = \mathbf{h}$ for all j . In this special case, the resulting SNR of the proposed multi-frame model is governed by both the power spectrum of the blur kernel and the spatial coverage of the crop operators. We define c_i to be the *coverage* at pixel i , i.e. the number of times pixel i is included in the windows $\{\mathbf{W}_1, \dots, \mathbf{W}_n\}$. The SNR for a multi-frame acquisition with coded illumination is bounded as

$$\text{SNR} \geq \sqrt{\min_i c_i} \cdot \frac{\gamma \bar{s}_0}{\text{DNF} \cdot \sqrt{\gamma \bar{s}_0 + \sigma_r^2}}. \quad (7)$$

Thus, the SNR for the multi-frame case is at least a factor of the square root of the minimum coverage c_i better than that of the single-frame case (Eq. (6)). The derivation is in Appendix 5.2.

Notably, the bound in Eq. (7) decouples the spatial coverage, determined by c_i , from the spectral quality of the blur, determined by DNF. This allows the decoupling of the motion path design from the illumination optimization. A good motion path ensures even spatial coverage through c_i , while a good illumination sequence ensures spectral coverage similar to single-frame methods. We focus system design on the maximization of this decoupled lower bound.

The decision to use the same blur kernel for each camera frame has several practical implications: the single illumination pattern is easy to store on a micro-controller with limited memory, and distorting all measurements by the same blurring operator makes post-processing registration simple. We additionally note that requiring the blurring motion to be along a single axis is not limiting, since in practice horizontal strips are reconstructed independently to accommodate computer memory. A limitation of this simplification is that our kernels are not optimized for nonlinear motion (such as around corners). We do not focus on this case, as it would provide only a marginal increase in field-of-view.

2.4. Illumination optimization

Previous work [32,34] showed that reconstructions performed using constant (non-coded) illumination will have very poor quality (in terms of SNR) compared to using optimized pulse

sequences or a short, single pulse. Here, we explore several approaches for generating illumination pulse sequences which maximize the reconstruction SNR (Eq. (6)). We first consider the problem of minimizing the DNF with respect to the kernel \mathbf{h} :

$$\begin{aligned} \text{DNF}(\gamma) &:= \min_{\mathbf{h}} \text{DNF} \\ \text{s.t. } & 0 \leq h_i \leq 1 \forall i, \quad \sum_i h_i = \gamma, \end{aligned} \quad (8)$$

where the inequality constraint on \mathbf{h} represents the finite optical throughput of the system. This optimization problem is non-convex, similar to previous work; our multiplexing factor γ is related to the kernel length N and throughput coefficient β in [32–34] by $\gamma = N\beta$. This definition enables a layered approach to maximizing the SNR: after solving Eq. (8) for each multiplexing factor, we find the one which optimizes Eq. (6) in the context of camera noise.

To simplify the optimization task, the positions encoded in the kernel \mathbf{h} may be restricted a priori, e.g. to a centered horizontal line with fixed length, as in Fig. 2. In the following, we constrain the positions to a straight line with length $N = 2\gamma$. This provides a sufficiently large sample space for kernel optimization, and is supported as optimal by analysis in [34].

We consider several methods for optimizing Eq. (8): random search over grayscale kernels, random search over binary kernels, and a projected gradient descent (PGD) approach. Our random search generates a fixed number of candidate kernels and chooses the one with the lowest DNF. The grayscale candidates were generated by sampling uniform random variables, while the binary were generated by sampling indices without replacement.

In our PGD approach, the kernel optimization problem in Eq. (8) is reformulated as the minimization of a smooth objective $g(\mathbf{h})$ subject to convex constraints \mathcal{S} . Starting from an initial kernel \mathbf{h}_0 , the update rule includes a gradient step followed by a projection:

$$\mathbf{h}^{k+1} = \text{Proj}_{\mathcal{S}}(\mathbf{h}^k - \alpha^k \nabla g(\mathbf{h}^k)). \quad (9)$$

Details of the reformulation and optimization approach are in Appendix 5.3.

Figure 3(a) shows the distribution of optimization results for 100 trials of each approach, where the random search methods sample 1000 candidates and PGD uses a step-size determined by backtracking line search until convergence from a random binary initialization. Example illumination sequences from each method and their corresponding power spectra are displayed in Fig. 3(b). Power spectra with low magnitudes correspond to a large DNF.

PGD and random binary search resulted in significantly better (lower) DNF than grayscale random search. Though the kernels with the lowest DNF were generated through PGD, binary random search results in comparable values and is up to 20× faster than PGD. A binary restriction also achieves fast illumination updates, since grayscale illumination (as in [33]) would require pulse-width-modulation spread across multiple clock cycles.

Plotting the DNF generated through binary random search for increasing multiplexing factor (Fig. 3(c)) reveals a concave curve. Fitting the curve with a power-law, a closed-form approximation for the DNF is $\text{DNF}(\gamma) = 1.12\gamma^{0.64}$. This analytic relationship allows for a direct optimization of the SNR: substituting any $\text{DNF}(\gamma) \propto \gamma^p$ into Eq. (7) and differentiating with respect to γ , we can determine the (approximately) optimal multiplexing factor γ^* as a function of mean strobed signal \bar{s}_0 and camera readout noise σ_r^2 for $p > 0.5$:

$$\gamma^* = \frac{2 - 2p}{2p - 1} \cdot \frac{\sigma_r^2}{\bar{s}_0}. \quad (10)$$

For smaller power law p (i.e. slower DNF growth with γ), the optimal multiplexing factor will be larger. When $p = 0.5$, the expression for SNR in Eq. (6) only increases with increasing multiplexing factors, meaning that the optimal multiplexing factor γ^* would be as large as possible

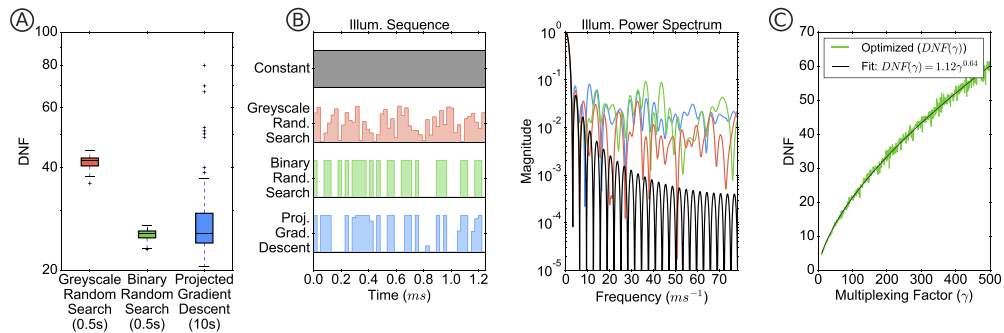


Fig. 3. A) Deconvolution noise factor (DNF) results for designing the illumination codes using different optimization methods. Random search over binary patterns gives similar performance as projected gradient descent (PGD), but is much faster. B) The optimized illumination sequences and their power spectra. C) Optimized DNF (binary random search method) for different values of the multiplexing factor γ . A power law fit of $DNF(\gamma) = 1.12\gamma^{0.641}$ has standard error of 1.075.

given hardware constraints. We show in Appendix 5.4 that $p = 0.5$ represents a lower bound on the size of the DNF, i.e. that $DNF(\gamma) \geq \gamma^{0.5}$ regardless of optimization method or illumination sequence. The experimental $p = 0.64$ accurately reflects the practical relationship between DNF and multiplexing factor. While part of that relationship may come from the increasing difficulty of optimization as the decision space grows with γ , this reflects actual limitations in practice.

The expression for optimal γ^* also solidifies the intuition that a larger multiplexing factor should be used for systems with high noise in order to increase detection SNR, while a lower multiplexing factor is appropriate for less noisy systems. This result is in agreement with [34], which demonstrated empirically that the best choice of multiplexing factor depends on the relative magnitude of the acquisition noise.

2.5. Experimental setup

Our system is built around an inverted microscope (TE300 Nikon) using a lateral motion stage (Prior, H117), as shown in Fig. 1. Images were acquired using a sCMOS camera (PCO.edge 5.5, PCO) through hardware triggering, and illumination was provided by a high-power LED (M470L3, Thorlabs) which was controlled by a micro-controller (Teensy 3.2, PJRC). Brightfield measurements were illuminated using one of two sources: a custom LED illuminator with 40 blue-phosphor LEDs (VAOL-3LWY4, VCC), or a single, high-power LED source (Thorlabs M470L3), both modulated using a simple single-transistor circuit through the same micro-controller. The first illuminator was designed to have a broad spectrum for brightfield imaging, while the second was intended for fluorescence imaging, having a narrow spectral bandwidth. For this project, we adopt very simple LED circuitry to avoid electronic speed limitations associated with dimming (due to pulse-width-modulation) and serial control of LED driver chips.

The micro-controller firmware was developed as part of a broader open-source LED array firmware project [37] and images were captured through the python bindings of Micro-Manager [38], which were controlled through a Jupyter notebook [39], enabling fast prototyping of both acquisition and reconstruction pipelines in the same application. With the exception of our custom illumination device, everything in our optical system is commercially available. Therefore, our method is amenable to most traditional optical setups and can be implemented through the simple addition of our temporally-coded light source and open-source software.

All acquisitions were performed using a 10 \times , 0.25NA objective. Because the depth-of-field (DOF) was relatively large (8.5 μm) compared to our sample, we were able to level the sample

manually prior to acquisition using adjustment screws on the motion stage, to ensure the sample remained in focus across an approximately 20mm movement range. For larger areas or shallower DOF, we anticipate a more sophisticated leveling technique will be necessary, such as active autofocus [40,41] or illumination-based autofocus [42].

3. Results

3.1. Gigapixel reconstruction

We first demonstrate a 1.18 gigapixel reconstruction of a microscope slide plated with $4.7\mu\text{m}$ polystyrene fluorescent beads (Thermo-Fisher) in Fig. 4. Each scan path consisted of up to 36 1D continuous scans structured in a raster-scanning pattern to enable fast 2D scanning of the sample. The number of scans was determined by the sample coverage on the slide; for arbitrarily large samples the limiting factor would be the range of the motion stage. For our system, this limit was approximately $114\times 75\text{mm}$.




Fig. 4. 1.18 Gigapixel $23\times 20\text{mm}$ full-field reconstruction of $4.7\mu\text{m}$ fluorescent microspheres. While stop-and-stare (S&S) measurements have the highest SNR, our coded illumination measurements were more than $5.5\times$ faster, while maintaining enough signal to distinguish individual microspheres, in contrast to strobed illumination, which is fast but noisy. Inset scale bars are $50\mu\text{m}$.

For comparison, we acquired stop-and-stare, strobed and coded illumination datasets and performed image stitching and registration of the three datasets for a direct comparison of image quality. The total acquisition time for the strobed and coded reconstructions of this size was 31.6 seconds, while a comparable stop-and-stare acquisition required 210.9 seconds. The computation time for coded-illumination reconstructions with step-size $\alpha = 0.5$ was approximately 30 minutes on a Macbook Pro (Apple) with attached RX580 external GPU (Advanced Micro-Devices), or approximately 2 minutes when parallelized across 18 EC2 p2.xlarge instances with Nvidia Titan GPUs (Amazon Web Services), excluding data transfer to and from our local machine.

3.2. Acquisition method comparison

To compare our coded illumination method with existing high-throughput imaging techniques, we quantify the expected SNR for each method based on relevant system parameters such as source illuminance, camera noise, and desired acquisition frame-rate. In conventional high-throughput imaging, the stop-and-stare strategy will provide higher SNR than strobed illumination, but is only feasible for low-frame rates due to mechanical limitations of the motion stage. Therefore,

we restrict our comparison to strobed illumination and coded illumination. In the next section, we perform a comprehensive analysis of these trade-offs.

We compare acquisitions with both brightfield and fluorescence configurations, while sweeping the output power of the LED source. We expect measurements acquired in a fluorescence configuration will have significantly lower signal due to the conversion efficiency of fluorophores. In each case, we measured the illuminance at the camera plane using an optical power meter (Thorlabs). In the fluorescence configuration, both the emission and excitation filters were in-place at the time of measurement. For each image we computed the average imaging SNR across three different regions of interest to produce an experimental estimate of the overall imaging SNR. To ensure a fair comparison, we performed no preprocessing on the data except for subtracting a known, constant background offset from each measurement which was characterized before acquisitions were performed and verified with the camera datasheet. Reconstructions for coded illumination were performed using the reconstruction algorithm described in Section 2.2.

Figure 5 shows that coded illumination can provide up to 10× higher SNR in low-illumination situations. Given typical hardware and noise parameters, coded illumination will be most useful for fluorescence microscopy (illuminance < 1000 lux). For brightfield microscopy at typical scan speeds, strobed illumination provides a higher SNR. The solid lines in Fig. 5 are theoretical predictions of reconstruction SNR for each method (described below) based on our system parameters, which are generally in agreement with experimental data.




Fig. 5. (A) Experimental SNR values for a USAF target imaged with varying illumination strength under strobed (squares) and coded illumination (diamonds). Solid lines show predicted SNR based on known system parameters. Experimental SNR values are the average of three SNR measurements performed across the field. Green and orange points represent inset data for fluorescent beads and resolution targets, respectively. Characteristic illuminance values are labeled for LED sources, Halogen-Tungsten Lamps (Hal), Mercury Lamps (Hg), Xenon Lamps (Xe), and Metal-Halide Lamps (M-H) [43]. (B) Example reconstructions and measurements with SNR and SBR values, highlighted in green if preferable or red if suboptimal compared to other methods.

In practice, users may incorporate regularization to improve reconstruction results by enforcing priors about the object. Figure 6 investigates the possible gains to be made by incorporating different regularization and denoising techniques. In this case, median filtering was the most effective technique for removing salt-and-pepper noise without introducing large artifacts. The displayed L2 regularization results are qualitatively similar to truncated Singular-Value Decomposition (TSVD), Total-Variation (TV), and Haar wavelet methods (not shown).

3.3. System component analysis

While the choice to use strobed or coded illumination depends largely on the illumination power of the source, other system parameters also affect this trade-off. Here, we consider camera noise and motion velocity, and include an analysis of when stop-and-stare should be used as opposed




Fig. 6. Reconstructions with different regularization methods: unregularized, L2 (coefficient 0.01) and median filtering (coefficient of 0.01, implemented by regularization by denoising [44]). While regularization decreases salt and pepper noise (see zoom-ins in bottom row), artifacts increase and are noticeable at larger spatial scales (top row).

to continuous-scan methods (strobed or coded illumination). As a first step, we derive the expected photons per pixel per second (J) that we expect to measure in a transmission microscope, incorporating system magnification (M), numerical aperture (NA), camera pixel size (Δ), mean wavelength ($\bar{\lambda}$), and the photometric look-up table $K(\bar{\lambda})$:

$$J = K(\bar{\lambda})\bar{\lambda}hc \cdot I_{lux} * (NA)^2 * \left(\frac{\Delta}{M}\right)^2, \quad (11)$$

where h is Planck's constant, c is the speed of light, I_{lux} is the source illuminance in lux, and J is the photon flux per pixel-second. Given J , the mean signal \bar{s} is a function of the illumination time t_{illum} and the camera quantum efficiency Q :

$$\bar{s} = JQt_{illum}. \quad (12)$$

Substituting Eq. (12) into Eq. (6), we define the expected SNR (used in Fig. 5 and Fig. 7) as a function of these parameters as well as the blur kernel h and camera readout noise σ_r :

$$SNR = \frac{JQt_{illum}}{DNF\sqrt{JQt_{illum} + \sigma_r^2}}. \quad (13)$$

The minimum pulse duration, t_{illum} , and DNF are functions which change based on acquisition strategy. For stop-and-stare and strobed acquisitions, we set $DNF = 1$, since no deconvolution is being performed, while for coded acquisitions DNF depends on the parameter γ as derived in Section 2.4. Similarly, t_{illum} is set based on acquisition strategy and motion stage parameters. For stop-and-stare illumination, t_{illum} is proportional to the residual time after stage movement, including stage velocity (v_{stage}), stage acceleration (a_{stage}), the camera FOV along the blur axis (FOV), mechanical settle time (t_{stage}), and desired acquisition frame rate r_{frame} :

$$t_{illum} = \frac{1}{r_{frame}} - 2\frac{v_{stage}}{a_{stage}} - \frac{FOV - 0.5 * \frac{v_{stage}^2}{a_{stage}}}{v_{stage}}. \quad (14)$$

For strobed and coded illumination, the minimum pulse duration t_{illum} is set by the overlap fraction (O) between frames, the number of pixels along the blur direction (N_{px}), and the multiplexing by

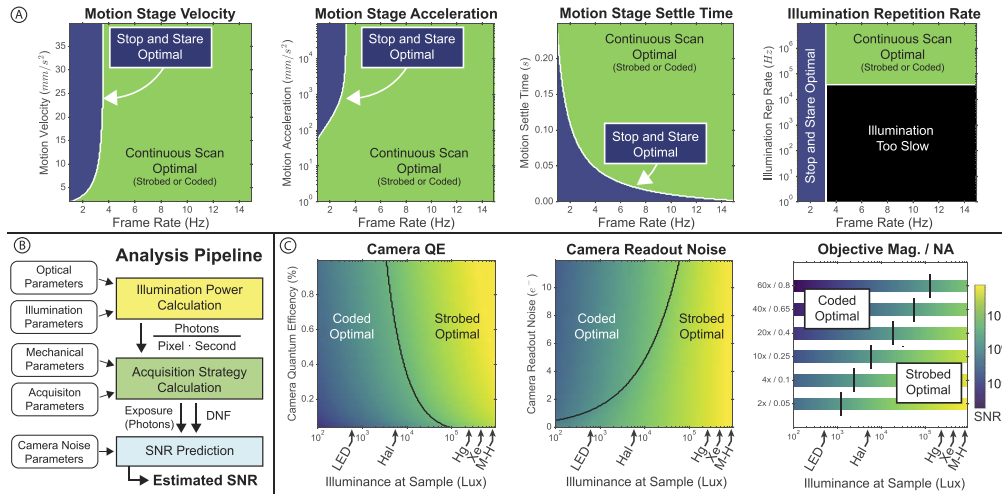


Fig. 7. Analysis of system parameters. A) The stop-and-stare acquisition strategy is optimal, but only possible for some configurations of mechanical system parameters and frame rates. B) Analysis pipeline for predicting SNR from system parameters, including illumination power, mechanical parameters of the motion stage, and camera noise parameters. C) Different combinations of optical system parameters and system illuminance determine the best possible SNR and whether strobed or coded illumination is preferable. Characteristic illuminance values for LED sources, Halogen-Tungsten Lamps (Hal), Mercury Lamps (Hg), Xenon Lamp (Xe), and Metal-Halide Lamps (M-H) are shown for reference.

γ (with $\gamma = 1$ for strobed illumination):

$$t_{illum} = \frac{\gamma}{r_{frame}(1 - O)N_{px}} \quad (15)$$

We implicitly calculate the velocity in terms of the fastest speed where two frames may overlap with O within a time set by r_{frame} . Derivations for the above relationships are in Appendix 5.5. With these theoretical values for t_{illum} , we derive closed-form solutions for SNR as a function of acquisition rate, with given system parameters (Table 1 in Appendix 5.6), using the pipeline in Fig. 7(b).

Our system analysis is divided into two parts: a mechanical comparison of stop-and-stare versus continuous motion, and an optical comparison between strobed and coded illumination. When stop-and-stare is possible given a desired acquisition frame rate, it will always provide higher SNR than strobed or coded illumination due to high photon counts and no deconvolution noise. Figure 7(a) analyzes where stop-and-stare is both possible and optimal compared to a continuous acquisition technique, as a function of frame rate. If the frame rate is low or limited by other factors (such as sample stability), stop-and-stare will be optimal in terms of SNR. If a high-frame rate is desired, however, a continuous acquisition strategy is optimal, so long as the illumination repetition rate of the source is fast enough to accommodate the sample speed.

Figure 7(c) describes the optimal continuous imaging technique as a function of illuminance, imaging objective, camera readout noise (σ_r), and camera quantum efficiency (QE). Generally speaking, higher illuminance values favor strobed illumination (being shot-noise limited), while lower illuminance values favor coded illumination (being read-noise limited). Conversely, as read noise (σ_r) increases or camera QE decreases, coded illumination becomes more beneficial. Practically, a camera with high read-noise and low QE will favor coded illumination more strongly (at higher source illuminance) than a high-end camera (such as the PCO.edge 5.5 used in this study), where $\sigma_r \approx 3.7e^-$ and $QE \approx 0.6$. In addition, objectives with a higher magnification and

NA will generally favor coded illumination more strongly due to the decreasing $\frac{NA}{Mag}$ ratio, which reduces the amount of light collected by the objective. It should be noted, however, that higher NA values will require more sophisticated autofocus methods than those presented in this work. Example illuminance values for common microscope sources were calculated based on estimated source power at 550nm [43].

3.4. Biological limitations for fluorescence microscopy

In fluorescence imaging, autofluorescence [45] and photobleaching [46] are primary considerations when assessing system throughput. Photobleaching, the result of chemical interactions of activated fluorophores with the surrounding medium, is of particular concern for motion deblurring applications due to a potential nonlinear response when a large number of illumination pulses are used. Practically, photobleaching can limit the maximum number of pulses which a sample may tolerate before exhibiting a non-linear response, causing strobed illumination to become a more favorable option, even at low-light. However, the region where this condition occurs is small (for the proposed system), and near the photobleaching limit (Fig. 8(a)). Autofluorescence is also a well-studied process which can further degrade the quality of fluorescence images. Autofluorescence affects contrast, which is best quantified using the contrast-to-noise ratio (CNR):

$$CNR = \frac{\gamma\bar{s}_0 - \bar{b}}{DNF\sqrt{\gamma\bar{s}_0 + \bar{b} + \sigma_r^2}}, \quad (16)$$

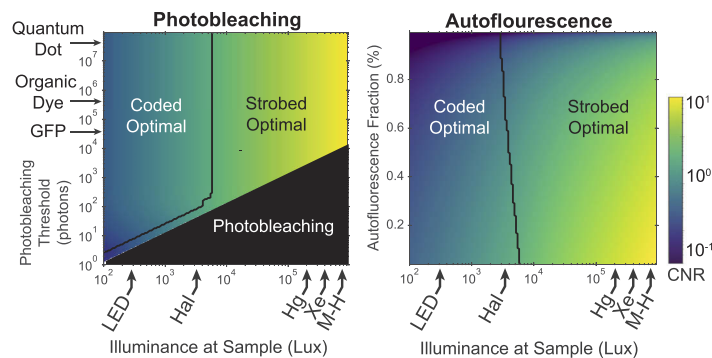


Fig. 8. Analysis of constraints imposed by the chemical fluorescence process, using contrast-to-noise ratio (CNR) as the figure of merit. Photobleaching influences the choice between coded and strobed illumination only when introducing a coding scheme would cause photobleaching, corresponding to a thin area of strobed optimality near the photobleaching limit. This plot assumes no background autofluorescence, so CNR and SNR are equivalent. The amount of autofluorescence relative to the signal mean has a slight effect on the optimality of strobed and coded illumination, but the effect is not strong relative to the other parameters studied here. Generally, the presence of autofluorescence degrades CNR ratio for all methods and illumination levels.

where \bar{b} is the mean background signal (autofluorescence). Note that in the absence of a background ($\bar{b} = 0$), CNR is equivalent to SNR. Figure 8(b) illustrates the relative optimality of strobed and coded illumination in the presence of background autofluorescence, expressed as a fraction of the primary signal.

The lifetime of various fluorophores is not limiting at illumination speeds presented in this work. Most endogenous fluorophores and fluorescent proteins have a lifetime of less than 10ns, while organic dyes may have lifetimes of less than 100ns [47]. The fastest illumination source used in this work had a repetition period of approximately 4 μ s, which is 40 \times faster than the

fluorophore-limited update rate for organic dyes. Still, for motion stages moving at high velocities (greater than $100 \frac{mm}{s}$) and high magnifications (greater than $40\times$), it will become more important to consider the lifetime of the dyes used. These same constraints would also apply to strobed imaging, but not stop-and-stare (which does not require high-speed signal modulation).

4. Conclusion

We have demonstrated a high-throughput imaging framework which employs multi-frame motion deblurring using temporally-coded illumination. Through both experiment and theoretical analysis we have shown the applicability of our method for fluorescence microscopy, and performed a comprehensive analysis of when our method is optimal in terms of source power and other system parameters. These results indicate that coded illumination provides up to $10\times$ higher SNR than conventional strobed illumination methods in low-light situations, making our method particularly well-suited for applications in drug-discovery and whole-slide imaging. Our analysis of optimal kernel selection indicates that efficient illumination sequences can be calculated quickly and efficiently using a random search, and our analysis of optimal pulse length provides an approximate relationship between the pulse sequence length and source illuminance. Further, our proposed multi-frame reconstruction algorithm produces high-quality results using simple accelerated gradient descent with no regularization, and can be scaled to multiple cloud instances for fast data processing. Future work should address more complicated motion pathways, self-calibration, and reconstructions using under-sampled data.

5. Open source

The software used in this study is available under an open-source (BSD 3-clause) license under two packages:

- Reconstruction and Acquisition Code [48]
- Micro-controller Firmware [37]

Appendix

5.1. Derivation of SNR

In this section we derive the expression for the SNR of a recovered image. Considering the additive noise acquisition model, $\mathbf{y} = \mathbf{A}\mathbf{x} + \eta$, the recovered image $\hat{\mathbf{x}}$ is given by:

$$\hat{\mathbf{x}} = \mathbf{A}^\dagger \mathbf{y} = \mathbf{x} + \mathbf{A}^\dagger \eta .$$

In what follows, we assume only that η is zero mean with covariance $\sigma_\eta^2 \mathbf{I}$. At the end of this section, we briefly outline how to extend this style of analysis to spatially-correlated noise or simple forms of regularization.

Defining the mean of the recovered object $\mu = \mathbb{E}[\hat{\mathbf{x}}]$, as well as the covariance $\Sigma = \mathbb{E}[(\hat{\mathbf{x}} - \mathbf{x})(\hat{\mathbf{x}} - \mathbf{x})^H]$, we calculate the imaging SNR using the root mean squared error (RMSE):

$$SNR = \frac{\frac{1}{m} \sum_{i=1}^m \mu_i}{\sqrt{\frac{1}{m} \text{Tr}(\Sigma)}} .$$

Assuming zero-mean noise, the numerator is the average object signal \bar{s} . Expanding the covariance term in the denominator,

$$\Sigma = \mathbb{E}[\mathbf{A}^\dagger \eta (\mathbf{A}^\dagger \eta)^H] = \sigma_\eta^2 \mathbf{A}^\dagger (\mathbf{A}^\dagger)^H ,$$

where we assume that the covariance of η is $\sigma_\eta^2 \mathbf{I}$. Then,

$$\text{Tr}(\mathbf{A}^\dagger (\mathbf{A}^\dagger)^\text{H}) = \sum_{i=1}^m \frac{1}{\sigma_i(\mathbf{A})^2} = \frac{1}{\sigma_1(\mathbf{A})^2} \sum_{i=1}^m \frac{\sigma_1(\mathbf{A})^2}{\sigma_i(\mathbf{A})^2}.$$

Thus we have that

$$SNR = \frac{\bar{s}}{\frac{1}{\sigma_1(\mathbf{A})} \sqrt{\frac{1}{m} \sum_{i=1}^m \frac{\sigma_1(\mathbf{A})}{\sigma_i(\mathbf{A})}} \cdot \sigma_\eta} := \frac{\sigma_1(\mathbf{A}) \bar{s}}{\text{DNF} \sigma_\eta},$$

where we use the general definition of the deconvolution noise factor (DNF). This expression is consistent with the definition in (5) for convolutional operators, where we note that the singular values are given by the power spectrum of the kernel \mathbf{h} . Further, we note that in this case $\sigma_1(\mathbf{A}) = \gamma$ since that is the DC component of a non-negative signal.

5.1.1. Extensions

Briefly, we outline how to extend the previous analysis in two cases: when the noise is spatially correlated and when simple regularization schemes are used.

First, we remark that if the noise η is spatially correlated, and instead has covariance Σ_η , the only change is a rescaling of the operator \mathbf{A} :

$$SNR = \frac{\bar{s} \sigma_1(\Sigma_\eta^{-1/2} \mathbf{A})}{\sqrt{\frac{1}{m} \sum_{i=1}^m \frac{\sigma_1(\Sigma_\eta^{-1/2} \mathbf{A})}{\sigma_i(\Sigma_\eta^{-1/2} \mathbf{A})}}}$$

which occurs because the computation for covariance changes as follows: $\Sigma = \mathbb{E}[\mathbf{A}^\dagger \eta (\mathbf{A}^\dagger \eta)^\text{H}] = \mathbf{A}^\dagger \Sigma_\eta (\mathbf{A}^\dagger)^\text{H}$.

Second, we consider simple singular value-based regularization. We demonstrate the extension for truncated SVD (TSVD), and remark that a similar analysis can be performed for L2 regularization. The truncated SVD objective with parameter k is

$$O(\mathbf{x}) = \frac{1}{2} \|\mathbf{A}_k \mathbf{x} - \mathbf{y}\|_2^2,$$

where \mathbf{A}_k contains only the top k singular values of the matrix \mathbf{A} , i.e.

$$\mathbf{A}_k = \mathbf{U} \text{diag}(\sigma_1(\mathbf{A}), \dots, \sigma_k(\mathbf{A}), 0, \dots, 0) \mathbf{V}^\text{H}.$$

We remark that this objective is very easy to implement for convolutional operators, whose singular values are given by the magnitude of the Fourier transform of the convolutional kernel. This objective introduces *bias* into the recovered image $\hat{\mathbf{x}}$:

$$\hat{\mathbf{x}} = \mathbf{A}_k^\dagger \mathbf{A}_k \mathbf{x} + \mathbf{A}_k^\dagger \eta = \mathbf{V}_k \mathbf{V}_k^\text{H} \mathbf{x} + \mathbf{A}_k^\dagger \eta.$$

Therefore, the mean and variance of the reconstructed image depend on characteristics of the true object. We define the following quantities, noting that various statistical assumptions about the content of the true object can result in further simplification. First, we denote as α_k the fraction of the mean signal that resides along the top k singular directions, i.e. within $\mathbf{V}_k \mathbf{V}_k^\text{H} \mathbf{x}$. Second, we denote as β_k the fraction of the signal variance that lies along the top k singular directions.

Then, the SNR expression is

$$SNR = \frac{\sigma_1(\mathbf{A})\bar{s}}{DNF_k\sigma_\eta},$$

where we define a bias-DNF DNF_k which depends on the underlying true object and the noise variance in addition to the operator \mathbf{A} :

$$DNF_k = \frac{1}{\alpha_k} \sqrt{(1 - \beta_k) \frac{\sigma_1(\mathbf{A})^2 \sigma_x^2}{\sigma_\eta^2} + \frac{1}{k} \sum_{i=1}^k \frac{\sigma_i(\mathbf{A})}{\sigma_i(\mathbf{A})}}.$$

Notice that since the sum is over the k largest singular values, depending on the content of the true object, DNF_k may be smaller than DNF and thus preferable.

5.2. Multi-frame decomposition

We consider the case of a multiframe operator with the same blur kernel \mathbf{h} used in every frame. In this case, the forward operator has the form

$$\mathbf{A} = \begin{bmatrix} \mathbf{W}_1 \\ \vdots \\ \mathbf{W}_n \end{bmatrix} \mathbf{B} := \mathbf{W}\mathbf{B}.$$

Following the derivation of SNR from the previous section, we compute $\text{Tr}(\mathbf{A}^\dagger(\mathbf{A}^\dagger)^\text{H})$. First,

$$\mathbf{A}^\dagger = (\mathbf{B}^\text{H}\mathbf{W}^\text{H}\mathbf{W}\mathbf{B})^{-1}\mathbf{B}^\text{H}\mathbf{W}^\text{H} = \mathbf{B}^{-1}(\mathbf{W}^\text{H}\mathbf{W})^{-1}\mathbf{W}^\text{H},$$

assuming that \mathbf{B} and $\mathbf{W}^\text{H}\mathbf{W}$ are invertable. Then we have

$$\begin{aligned} \text{Tr}(\mathbf{A}^\dagger(\mathbf{A}^\dagger)^\text{H}) &= \text{Tr}(\mathbf{B}^{-1}(\mathbf{W}^\text{H}\mathbf{W})^{-1}\mathbf{W}^\text{H}\mathbf{W}(\mathbf{W}^\text{H}\mathbf{W})^{-\text{H}}\mathbf{B}^{-\text{H}}) \\ &= \text{Tr}(\mathbf{B}^{-1}(\mathbf{W}^\text{H}\mathbf{W})^{-1}\mathbf{B}^{-1}) = \text{Tr}((\mathbf{W}^\text{H}\mathbf{W})^{-1}\mathbf{B}^{-2}). \end{aligned}$$

We now consider the form of $\mathbf{W}^\text{H}\mathbf{W} = \sum_{j=1}^n \mathbf{W}_j^\text{H}\mathbf{W}_j$. Each $\mathbf{W}_j^\text{H}\mathbf{W}_j$ is a square diagonal matrix with either a 0 or 1 for each diagonal entry, depending on whether the corresponding pixel is included in the window. Thus the sum $\mathbf{W}^\text{H}\mathbf{W}$ is a diagonal matrix with the i th diagonal value given by the number of times pixel i is included in the windows $\{\mathbf{W}_1, \dots, \mathbf{W}_n\}$, a quantity we denote as $c_i = \sum_{j=1}^n \mathbf{W}_j \mathbf{e}_i$ where $\{\mathbf{e}_i\}$ are the standard basis vectors.

Before we proceed further, note that for any matrices \mathbf{M} and \mathbf{D} with non-negative entries and \mathbf{D} diagonal,

$$\text{Tr}(\mathbf{D}\mathbf{M}) = \sum_i D_{ii}M_{ii} \leq \max_i D_{ii} \cdot \text{Tr}(\mathbf{M}).$$

We can therefore conclude that

$$\text{Tr}(\mathbf{A}^\dagger(\mathbf{A}^\dagger)^\text{H}) \leq \max_i \frac{1}{c_i} \cdot \text{Tr}(\mathbf{B}^{-2}) = \frac{1}{\min_i c_i} \cdot \sum_{i=1}^m \frac{1}{|\tilde{\mathbf{h}}_i|^2}.$$

Thus we see that the expression for the covariance is decreased by a factor of at least the square root of minimum coverage. This corresponds to the lower bound on the SNR:

$$SNR \geq \sqrt{\min_i c_i} \cdot \frac{\gamma\bar{s}}{DNF\sigma_\eta},$$

where DNF is defined as in (5).

5.3. Blur kernel optimization

In this section we discuss the reformulation of the optimization problem in (10) as a smooth objective with convex constraints. Recall that the optimization problem has the form

$$\min_{\mathbf{h}} \sqrt{\frac{1}{m} \sum_{i=0}^m \frac{\max_i |\tilde{\mathbf{h}}|_i^2}{|\tilde{\mathbf{h}}|_i^2}} \quad s.t. \quad 0 \leq h_i \leq 1 \quad \forall i, \quad \sum_i h_i = \gamma.$$

First, note that by definition $\tilde{\mathbf{h}} = \mathbf{F}\mathbf{h}$ where \mathbf{F} represents the discrete Fourier transform (DFT) matrix. Then, we know that $\max_i |\tilde{\mathbf{h}}|_i^2$ is the DC component of the signal, which is equal to γ and therefore fixed for any feasible \mathbf{h} . Therefore, the blur kernel which maximizes (8) is the same as the one that maximizes

$$\min_{\mathbf{h}} \sum_{i=0}^m \frac{1}{(\mathbf{F}_i^H \mathbf{h})^2} \quad s.t. \quad 0 \leq h_i \leq 1, \quad \sum_i h_i = \gamma,$$

where \mathbf{F}_i represents columns of the DFT matrix.

It is possible to use projected gradient methods because the objective function is smooth nearly everywhere and the constraints are convex. At each iteration, there is a gradient step followed by a projection step. The gradient step is defined as

$$\tilde{\mathbf{h}}^{k+1} = \mathbf{h}^k + \alpha^k \sum_{i=0}^m \frac{2}{(\mathbf{F}_i^H \mathbf{h}^k)^3} \cdot \mathbf{F}_i,$$

for potentially changing step size α^k . The projection step is defined as

$$\mathbf{h}^{k+1} = \Pi_{\mathcal{S}}(\tilde{\mathbf{h}}_{k+1}),$$

where \mathcal{S} is the intersection of the box constraint $\{0 \leq h_i \leq 1\}$ and the simplex constraint $\{\sum_i h_i = \gamma\}$. Efficient methods for this projection exist [49].

5.4. Fundamental DNF limits

There are fundamental limits on how SNR can be improved by coded illumination. We examine a fundamental lower bound on the DNF to demonstrate this.

Recall that

$$\text{DNF}^2 = \max_i |\tilde{\mathbf{h}}|_i^2 \cdot \frac{1}{m} \sum_{i=1}^m \frac{1}{|\tilde{\mathbf{h}}|_i^2}.$$

Then, note that $\frac{1}{m} \sum_{i=1}^m \frac{1}{|\tilde{\mathbf{h}}|_i^2}$ is the reciprocal of the harmonic mean of $\{|\tilde{\mathbf{h}}|_1^2, \dots, |\tilde{\mathbf{h}}|_m^2\}$. Since the harmonic mean is always less than the arithmetic mean, we have that

$$\frac{1}{m} \sum_{i=1}^m \frac{1}{|\tilde{\mathbf{h}}|_i^2} \geq \frac{1}{\frac{1}{m} \sum_{i=1}^m |\tilde{\mathbf{h}}|_i^2}.$$

Next, we apply Parseval's theorem and have $\frac{1}{m} \sum_{i=1}^m |\tilde{\mathbf{h}}|_i^2 = \sum_{i=1}^m h_i^2$. Additionally, $\max_i |\tilde{\mathbf{h}}|_i$ is the DC component of the signal, which is specified by the constraint $\sum_{i=1}^m h_i = \gamma$. As a result,

$$\text{DNF}^2 \geq \frac{\gamma^2}{\sum_{i=1}^m h_i^2}.$$

Finally, we see that

$$\max_{\mathbf{h} \in [0,1]^m} \sum_{i=1}^m h_i^2 : \sum_{i=1}^m h_i = \gamma$$

is achieved for binary h and has the maximum value γ . Thus,

$$\text{DNF}^2 \geq \frac{\gamma^2}{\sum_{i=1}^m h_i^2} \geq \frac{\gamma^2}{\gamma} = \gamma.$$

That is, the DNF grows at a rate of at least $\sqrt{\gamma}$. As a result, the best achievable SNR (using (6)) is

$$\text{SNR} \leq \frac{\sqrt{\gamma \bar{s}_0}}{\sqrt{\gamma \bar{s}_0 + \sigma_r^2}} = \sqrt{\bar{s}_0} \sqrt{\frac{\gamma \bar{s}_0}{\gamma \bar{s}_0 + \sigma_r^2}},$$

This upper bound on SNR increases with γ . In Methods Section 2.4, we discuss an exact closed form for $\text{DNF}(\gamma)$ that yields an expression for optimal multiplexing.

However, if σ_r is much smaller than the total captured signal, i.e. $\sigma_r \ll \gamma \bar{s}_0$, the SNR will not increase with γ , and in fact its maximum value,

$$\text{SNR} \leq \sqrt{\bar{s}_0}$$

is achieved by strobed illumination (i.e. $\gamma = 1$). In other words, when signal is large compared with readout noise, strobed will be optimal, regardless of the illumination optimization method.

5.5. Derivation of illumination throughput

5.5.1. Stop-and-stare

In the stop-and-stare acquisition strategy, the sample is illuminated for the full dwell time (t_{sns}), which is set by motion stage parameters such as maximum velocity, acceleration, and the necessary stage settle time (v_{stage} , a_{stage} , and t_{settle} respectively), as well as camera readout (t_{readout}). These parameters are related to frame rate r_{frame} by the following relationship:

$$t_{\text{sns}} = \frac{1}{r_{\text{frame}}} - \max(t_{\text{readout}}, 2t_{\text{accel}} + t_{\text{move}})$$

Note that this equation assumes perfect hardware synchronization and instantaneous acceleration ($\frac{\partial a}{\partial t} = \infty$). The variables t_{accel} and t_{move} are defined as:

$$t_{\text{accel}} = \frac{v_{\text{stage}}}{a_{\text{stage}}}$$

$$t_{\text{move}} = \frac{d_{\text{frame}} - 0.5 * a_{\text{stage}} * t_{\text{accel}}^2}{v_{\text{stage}}}$$

Here the expression $d_{\text{frame}} = \text{FOV} * (1 - O)$ is the distance between frames, which is determined by the field-of-view of a single frame (FOV) and inter-frame overlap fraction O .

Combining terms, we arrive at an expression for t_{sns} :

$$t_{\text{sns}} = \frac{1}{r_{\text{frame}}} - \max\left(t_{\text{readout}}, 2 \frac{v_{\text{stage}}}{a_{\text{stage}}} - \frac{d_{\text{frame}} - 0.5 * a_{\text{stage}} * t_{\text{accel}}^2}{v_{\text{stage}}}\right)$$

When camera readout time t_{readout} is short, t_{sns} can be simplified to:

$$t_{\text{sns}} = \frac{1}{r_{\text{frame}}} - 2 \frac{v_{\text{stage}}}{a_{\text{stage}}} - \frac{d_{\text{frame}} - 0.5 * \frac{v_{\text{stage}}^2}{a_{\text{stage}}}}{v_{\text{stage}}}$$

5.5.2. Strobed illumination

The maximum pulse duration for strobed illumination is related to the time required to move a distance of one effective pixel size $\frac{\Delta}{M}$ at a velocity v_{stage} :

$$t_{strobe} = \frac{\frac{\Delta}{M}}{v_{stage}}$$

The stage velocity v_{stage} may be bounded by the motion stage hardware (v_{max}) or by the FOV of the microscope:

$$v_{stage} = \min(v_{max}, r_{frame}FOV)$$

5.5.3. Coded illumination

The calculation of t_{coded} for coded illumination is synonymous to the strobed illumination case, weighted by the multiplexing coefficient used to generate the illumination sequence (γ), and using the v_{stage} calculation from the strobed subsection:

$$t_{coded} = \frac{\gamma \frac{\Delta}{M}}{v_{stage}}$$

5.6. System parameters

Table 1. System parameters.

| Parameter | Value |
|-------------------------------|----------------------|
| Maximum Motion Stage Velocity | $40 \frac{m}{s}$ |
| Motion Stage Acceleration | $400 \frac{m}{s^2}$ |
| Motion Stage Settle Time | 0.1s |
| Objective Mag / NA | $10 \times / 0.25NA$ |
| Frame Overlap | 20% |
| Illumination Power | 600 |
| Camera Readout Time | 26ms |
| Camera Readout Noise | $3.7 e^-$ |
| Camera Quantum Efficiency | 60% |
| Camera Pixel Size | $6.5 \mu m$ |
| Fluorophore Quantum Yield | 79% lux |
| Illumination repetition Rate | 250kHz |

Funding

Qualcomm; Gordon and Betty Moore Foundation; David and Lucile Packard Foundation; National Science Foundation (DGE 1752814); Office of Naval Research (N00014-17-1-2191, N00014-17-1-2401, N00014-18-1-2833); Defense Advanced Research Projects Agency (FA8750-18-C-0101, W911NF-16-1-0552); Amazon Web Services; Chan Zuckerberg Initiative.

Acknowledgments

The authors would like to thank Li-Hao Yeh, Shwetadwip Chowdhury, Emrah Boston, and David Ren for useful discussions and assistance with sample preparation.

Disclosures

Z.P. and L.W. are co-founders of Spectral Coded Illumination Inc., which manufactures LED arrays similar to those presented in this work. The illumination devices shown here were not manufactured by Spectral Coded Illumination Inc.

References

1. Z. E. Perlman, M. D. Slack, Y. Feng, T. J. Mitchison, L. F. Wu, and S. J. Altschuler, "Multidimensional drug profiling by automated microscopy," *Science* **306**(5699), 1194–1198 (2004).
2. P. Brodin and T. Christophe, "High-content screening in infectious diseases," *Curr. Opin. Chem. Biol.* **15**(4), 534–539 (2011).
3. M. Bickle, "The beautiful cell: high-content screening in drug discovery," *Anal. Bioanal. Chem.* **398**(1), 219–226 (2010).
4. U. Liebel, V. Starkuviene, H. Erfle, J. C. Simpson, A. Poustka, S. Wiemann, and R. Pepperkok, "A microscope-based screening platform for large-scale functional protein analysis in intact cells," *FEBS Lett.* **554**(3), 394–398 (2003).
5. W.-K. Huh, J. V. Falvo, L. C. Gerke, A. S. Carroll, R. W. Howson, J. S. Weissman, and E. K. O'shea, "Global analysis of protein localization in budding yeast," *Nature* **425**(6959), 686–691 (2003).
6. J. Peiffer, F. Majewski, H. Fischbach, J. Bierich, and B. Volk, "Alcohol embryo-and fetopathy: Neuropathology of 3 children and 3 fetuses," *J. Neurol. Sci.* **41**(2), 125–137 (1979).
7. M. Rimmelinck, M. B. S. Lopes, N. Nagy, S. Rorive, K. Rombaut, C. Decaestecker, R. Kiss, and I. Salmon, "How could static telepathology improve diagnosis in neuropathology?" *Anal. Cell. Pathol.* **21**(3-4), 177–182 (2000).
8. M. Alegro, P. Theofilas, A. Nguy, P. A. Castruita, W. Seeley, H. Heinsen, D. M. Ushizima, and L. T. Grinberg, "Automating cell detection and classification in human brain fluorescent microscopy images using dictionary learning and sparse coding," *J. Neurosci. Methods* **282**, 20–33 (2017).
9. C. Z. Microscopy, "Zeiss axio scan.z1," (2019).
10. G. Zheng, R. Horstmeyer, and C. Yang, "Wide-field, high-resolution Fourier Ptychographic microscopy," *Nat. Photonics* **7**(9), 739–745 (2013).
11. E. Betzig, G. H. Patterson, R. Sougrat, O. W. Lindwasser, S. Olenych, J. S. Bonifacino, M. W. Davidson, J. Lippincott-Schwartz, and H. F. Hess, "Imaging intracellular fluorescent proteins at nanometer resolution," *Science* **313**(5793), 1642–1645 (2006).
12. M. J. Rust, M. Bates, and X. Zhuang, "Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (storm)," *Nat. Methods* **3**(10), 793–796 (2006).
13. M. G. Gustafsson, "Surpassing the lateral resolution limit by a factor of two using structured illumination microscopy," *J. Microsc.* **198**(2), 82–87 (2000).
14. J. M. Rodenburg and H. M. Faulkner, "A phase retrieval algorithm for shifting illumination," *Appl. Phys. Lett.* **85**(20), 4795–4797 (2004).
15. L. Tian, X. Li, K. Ramchandran, and L. Waller, "Multiplexed coded illumination for fourier ptychography with an led array microscope," *Biomed. Opt. Express* **5**(7), 2376–2389 (2014).
16. N. Farahani, A. V. Parwani, and L. Pantanowitz, "Whole slide imaging in pathology: advantages, limitations, and emerging perspectives," *Pathol. Lab. Med. Int.* **7**, 23–33 (2015).
17. A. W. Lohmann, R. G. Dorsch, D. Mendlovic, Z. Zalevsky, and C. Ferreira, "Space-bandwidth product of optical signals and systems," *J. Opt. Soc. Am. A* **13**(3), 470–473 (1996).
18. L. Tian, Z. Liu, L.-H. Yeh, M. Chen, J. Zhong, and L. Waller, "Computational illumination for high-speed in vitro fourier ptychographic microscopy," *Optica* **2**(10), 904–911 (2015).
19. H. Wang, Y. Rivenson, Y. Jin, Z. Wei, R. Gao, H. Günaydin, L. A. Bentolila, C. Kural, and A. Ozcan, "Deep learning enables cross-modality super-resolution in fluorescence microscopy," *Nat. Methods* **16**(1), 103–110 (2019).
20. T. Nguyen, Y. Xue, Y. Li, L. Tian, and G. Nehmetallah, "Deep learning approach for fourier ptychography microscopy," *Opt. Express* **26**(20), 26470–26484 (2018).
21. Y. Rivenson, Z. Göröcs, H. Günaydin, Y. Zhang, H. Wang, and A. Ozcan, "Deep learning microscopy," *Optica* **4**(11), 1437–1443 (2017).
22. Y. Xue, S. Cheng, Y. Li, and L. Tian, "Reliable deep-learning-based phase imaging with uncertainty quantification," *Optica* **6**(5), 618–629 (2019).
23. J. Ho, A. V. Parwani, D. M. Jukic, Y. Yagi, L. Anthony, and J. R. Gilbertson, "Use of whole slide imaging in surgical pathology quality assurance: design and pilot validation studies," *Hum. Pathol.* **37**(3), 322–331 (2006).
24. G. Lepage, J. Bogaerts, and G. Meynants, "Time-delay-integration architectures in cmos image sensors," *IEEE Trans. Electron Devices* **56**(11), 2524–2533 (2009).
25. L. T. Grinberg, R. E. de Lucena Ferretti, J. M. Farfel, R. Leite, C. A. Pasqualucci, S. Rosemberg, R. Nitrini, P. H. N. Saldiva, and W. Jacob Filho, Brazilian Aging Brain Study Group, "Brain bank of the brazilian aging brain study group—a milestone reached and more than 1,600 collected brains," *Cell Tissue Banking* **8**(2), 151–162 (2007).
26. Z. F. Phillips, M. V. D'Ambrosio, L. Tian, J. J. Rulison, H. S. Patel, N. Sadras, A. V. Gande, N. A. Switz, D. A. Fletcher, and L. Waller, "Multi-contrast imaging and digital refocusing on a mobile microscope with a domed led array," *PLoS One* **10**(5), e0124938 (2015).

27. Z. Phillips, R. Eckert, and L. Waller, "Quasi-dome: A self-calibrated high-na led illuminator for fourier ptychography," in *Imaging Systems and Applications* (Optical Society of America, 2017), pp. IW4E–5.
28. Y. E. Nesterov, "A method for solving the convex programming problem with convergence rate $o(1/k^2)$," *Dokl. Akad. Nauk SSSR* **269**, 543–547 (1983).
29. K. Mitra, O. S. Cossairt, and A. Veeraraghavan, "A framework for analysis of computational imaging systems: Role of signal prior, sensor noise and multiplexing," *IEEE Trans. Pattern Anal. Mach. Intell.* **36**(10), 1909–1921 (2014).
30. K. Hagiwara and K. Kuno, "Regularization learning and early stopping in linear networks," in *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. IJCNN 2000. Neural Computing: New Challenges and Perspectives for the New Millennium*, vol. 4 (IEEE, 2000), pp. 511–516.
31. P. Yalamanchili, U. Arshad, Z. Mohammed, P. Garigipati, P. Entschew, B. Kloppenborg, J. Malcolm, and J. Melonakos, "ArrayFire - A high performance software library for parallel computing with an easy-to-use API," (2015).
32. R. Raskar, A. Agrawal, and J. Tumblin, "Coded exposure photography: motion deblurring using fluttered shutter," *ACM Trans. Graph.* **25**(3), 795–804 (2006).
33. C. Ma, Z. Liu, L. Tian, Q. Dai, and L. Waller, "Motion deblurring with temporally coded illumination in an led array microscope," *Opt. Lett.* **40**(10), 2281–2284 (2015).
34. A. Agrawal and R. Raskar, "Optimal single image capture for motion deblurring," in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, (IEEE, 2009), pp. 2560–2567.
35. O. Cossairt, M. Gupta, and S. K. Nayar, "When does computational imaging improve performance?" *IEEE Trans. on Image Process.* **22**(2), 447–458 (2013).
36. *Bounds for the condition numbers of spatially-variant convolution matrices in image restoration problems* (OSA, Toronto, 2011).
37. Z. Phillips, "Illuminate: Open-source led array control for microscopy" <https://github.com/zfphil/illuminate> (2019).
38. A. Edelstein, N. Amodaj, K. Hoover, R. Vale, and N. Stuurman, "Computer control of microscopes using mu;manager," *Curr. Protoc. Mol. Biol.* **92**(1), 14.20.1–14.20.17 (2010).
39. T. Kluyver, B. Ragan-Kelley, F. Pérez, B. Granger, M. Bussonnier, J. Frederic, K. Kelley, J. Hamrick, J. Grout, S. Corlay, P. Ivanov, D. Avila, S. Abdalla, and C. Willing, "Jupyter notebooks – a publishing format for reproducible computational workflows," in *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, F. Loizides and B. Schmidt, eds. (IOS Press, 2016), pp. 87–90.
40. "Nikon perfect focus," <https://www.microscopyu.com/applications/live-cell-imaging/nikon-perfect-focus-system>. Accessed: 2019-01-23.
41. "Zeiss definite focus," <https://www.zeiss.com/microscopy/us/products/light-microscopes/axio-observer-for-biology/definite-focus.html>. Accessed: 2019-01-23.
42. H. Pinkard, Z. Phillips, A. Babakhani, D. A. Fletcher, and L. Waller, "Deep learning for single-shot autofocus microscopy," *Optica* **6**(6), 794–797 (2019).
43. C. S. Murphy and M. W. Davidson, "Light source power levels," <http://zeiss-campus.magnet.fsu.edu/articles/lightsources/powertable.html> (2019). Accessed: 2019-02-13.
44. Y. Romano, M. Elad, and P. Milanfar, "The little engine that could: Regularization by denoising (red)," *SIAM J. Imaging Sci.* **10**(4), 1804–1844 (2017).
45. A. C. Croce and G. Bottiroli, "Autofluorescence spectroscopy and imaging: a tool for biomedical research and diagnosis," *Eur. J. Histochem.* **58**(4), 2461 (2014).
46. J. Lippincott-Schwartz, N. Altan-Bonnet, and G. H. Patterson, "Photobleaching and photoactivation: following protein dynamics in living cells." *Nature cell biology* pp. S7–14 (2003).
47. M. Y. Berezin and S. Achilefu, "Fluorescence lifetime measurements and biological imaging," *Chem. Rev.* **110**(5), 2641–2684 (2010).
48. Z. Phillips and S. Dean, "htdeblur, open-source acquisition and processing code for high-throughput microscopy using motion deblurring." <https://github.com/zfphil/htdeblur> (2019).
49. M. D. Gupta, S. Kumar, and J. Xiao, "L1 projections with box constraints," arXiv preprint arXiv:1010.0141 (2010).