Beyond Nyquist sampling: a cost-based approach

Ayça Özçelikkale* and Haldun M. Ozaktas

Department of Electrical Engineering, Bilkent University, Ankara TR-06800, Turkey
*Corresponding author: ayca@ee.bilkent.edu.tr

A sampling-based framework for finding the optimal representation of a finite energy optical field using a finite number of bits is presented. For a given bit budget, we determine the optimum number and spacing of the samples in order to represent the field with as low error as possible. We present the associated performance bounds as trade-off curves between the error and the cost budget. In contrast to common practice, which often treats sampling and quantization separately, we explicitly focus on the interplay between limited spatial resolution and limited amplitude accuracy, such as whether it is better to take more samples with lower amplitude accuracy or fewer samples with higher accuracy. We illustrate that in certain cases sampling at rates different from the Nyquist rate is more efficient.   © 2013 Optical Society of America

OCIS codes: 030.0030, 070.2025, 070.0070, 110.3055, 100.3020, 350.5730.

1. INTRODUCTION

This article addresses the problem of efficient representation of a finite-energy nonstationary optical field using a finite number of bits. We consider the scenario where a finite number of equidistant samples of the field is used for the representation. Each sample is of finite accuracy; that is, there is a finite number of distinguishable amplitude levels in each sample. Therefore, one can use a finite number of bits to represent each sample. The total number of bits used for all of the samples constitutes the bit cost associated with the representation. For a given bit cost budget, we determine the optimum number and spacing of the samples in order to represent the field with as low error as possible. Our framework leads to performance bounds associated with the finite representation of optical fields, in the form of trade-off curves between the representation cost and the error. We investigate the optimal values of parameters, such as the spacing, number, and accuracy of the samples, in order to achieve the most advantageous cost-error trade-off. We also discuss how the degree of coherence of the field and the level of noise affects our results.

Our approach emphasizes an important and sometimes overlooked issue in representation of optical fields, namely the limited amplitude accuracy of the samples. When a signal is to be represented with its samples, the Shannon–Nyquist sampling theorem is often used as a guideline. The theorem states that a band-limited signal with maximum frequency $B/2$ Hz can be recovered from its equidistant samples taken $1/B$ apart. In practice, signals may not be exactly band-limited, but rather effectively band-limited in the sense that the signal energy beyond a certain frequency is negligible. In such cases, the effective bandwidth is often used to determine a sampling interval. Another practical constraint is the impossibility of taking an infinite number of samples. Thus, it is common to determine an effective spatial extent L in the sense that the signal energy is negligible outside this extent, and use only the samples that fall in this effective spatial extent. This approach leaves us with a finite number LB of

samples. This approach may not always be the most appropriate manner in which to use the Shannon–Nyquist sampling theorem; there may be cases where one can do better by incorporating other available information. In particular, consider the practical scenario where the field is to be represented with a finite number of finite accuracy samples. Use of the conventional approach in this scenario raises a number of issues. For one thing, the concept of effective bandwidth and effective spatial extent is intrinsically ambiguous, in that there is some arbitrariness in deciding beyond what point the signal may be assumed negligible. This approach also completely ignores the fact that the samples will have limited amplitude accuracy. When we are required to represent the signal with a prespecified number of bits, the sampling interval dictated by the conventional sampling theorem may not be optimal. For instance, depending on the circumstances, it may be preferable to work with a larger sampling interval and a higher number of amplitude levels. In order to find the optimal values of these parameters, we must abandon the conventional approach and jointly optimize over the sampling interval and amplitude accuracies. Even when the amplitude accuracies are so high that we can assume the sample values to be nearly exact, the conventional sampling theorem may still not predict the optimal sampling interval if we are required to represent the signal with a given finite number of samples (especially when that number is relatively small).

One of the questions we ask in this context is the following: given that in practice the samples will have limited amplitude accuracy, is it possible to achieve lower reconstruction errors by choosing to sample at a rate different than the Nyquist rate? Although one may expect to compensate for the limited accuracy of the samples by oversampling, the precise relationships between the sampling parameters and the reconstruction error are not immediately evident. In this paper we give quantitative answers to this question by determining the optimal sampling parameters and the resulting performance bounds for the best achievable error for a given bit budget.

A study of these issues can also lead to a better understanding of what happens to the information carried by a wave field as it propagates. How well the field values at particular points in space can represent the whole field, has to do with how much information these values carry about the rest of the field. Many works in different areas, such as optics, electromagnetics, and information theory, have studied different aspects of the information-theoretical relationships in propagating wave fields. The concept of number of degrees of freedom is used in several works including [1–18]. Other works have adopted a sampling theory approach [19–22]. The concepts of structural and metrical information discussed in [23] have found application in [2,24,25]. Some researchers have focused on computational issues, where the aim is to process the signals without losing any significant information, as well as by using as little computational resources as possible, such as [26–29]. A number of works utilizing information theoretic concepts such entropy or channel capacity in different contexts have appeared [30–46]. Our purpose is to contribute to the understanding of the information theoretical relationships in optical fields by focusing on the trade-offs between amplitude accuracy and spatial resolution, and the trade-offs between the error and the bit budget.

We now present a brief overview of our article. In Section 2, we present the problem formulation, where we consider a general framework where the samples are taken after the field passes through a linear system. In Section 3, we present Gaussian–Schell model (GSM), the random field model used in our numerical experiments. We show the invariance of our error-cost trade-off curves for GSM fields propagating through first-order systems in Section 4. In Section 5, we present the optimum sampling strategies and the error-cost trade-off curves. We compare our optimal trade-off curves with the ones that would be obtained if the Shannon–Nyquist sampling theorem was used as the guideline in Section 6. In Section 7, we provide a general discussion on representation of optical fields using finite numbers of bits. Some concluding remarks are given in Section 8.

2. PROBLEM FORMULATION

Let the input field $f(x)$ reside in the $z = 0$ plane, which is perpendicular to the optical axis z. Considering only one transverse dimension for simplicity, let $f(x)$ be a zero-mean finite-energy proper complex Gaussian random field (random process). $f(x)$ passes through a possibly noisy linear system to produce the output $g(x)$:

$$g(x) = \mathcal{L}\{f(x)\} + n(x), \tag{1}$$

where $\mathcal{L}\{.\}$ denotes the linear optical system, and $n(x)$ is a random field denoting the system noise. $n(x)$ is modeled as a zero-mean proper complex Gaussian random field. We assume that the unknown random field $f(x)$ and system noise $n(x)$ are statistically independent.

M finite-accuracy equidistant samples of $g(x)$ are taken with the sampling interval Δ_x. The limited amplitude accuracy of the samples is modeled through an additive noise field

$$s_i = g(\xi_i) + m_i, \tag{2}$$

where $x = \xi_1, ..., \xi_M \in \mathbb{R}$ are the equidistant sampling locations with the spacing Δ_x, and the midpoint $x_0 = 0.5(\xi_1 + \xi_M)$.

We assume that the m_i's are independent, zero-mean, proper complex Gaussian random variables. We further assume that the m_i's are statistically independent of $f(x)$ and $n(x)$. By putting s_i in vector form, we obtain $\mathbf{s} = [s_1, ..., s_M]^T$. Before moving on, we underline that m_i models the uncertainty introduced by the measurement process. We usually have some degree of control over it; for instance, we can reduce it by using more accurate devices, provided we are willing to pay the price. On the other hand, the system noise $n(x)$ corresponds to other sources of uncertainty that are independent of the measurement process, and that we mostly have no control over, such as the thermal noise in the environment, uncontrolled fluctuations in the light source, or another information bearing signal, which is of no interest to us.

There is a cost associated with each sample. The cost associated with the ith sample is given by $C_{s_i} = \log_2(\sigma_{s_i}^2/\sigma_{m_i}^2)$ and is measured in bits. Here, $\sigma_{s_i}^2 = E[|s_i|^2]$ and $\sigma_{m_i}^2 = E[|m_i|^2]$, so that $\sigma_{s_i}/\sigma_{m_i}$ is essentially the ratio of the spread of the signal to the spread of the uncertainty, which corresponds to the number of distinguishable levels (dynamic range). Hence, the logarithm of this number may be considered to provide a measure of the number of bits needed to represent this variable. For a field value at a given location, smaller noise levels (smaller $\sigma_{m_i}^2$) correspond to a sample with higher amplitude accuracy and higher cost. On the other hand, a larger noise level corresponds to lower amplitude accuracy and lower cost. Further discussion of this cost function and some applications can be found in [44,47,48]. Here we will assume that the accuracy (hence the cost) associated with each sample is the same; that is, $C_{s_i} = C_{s_1}$, $i = 1, ..., M$. The total cost of the representation is then simply $C_T = \sum_{i=1}^M C_{s_i} = M C_{s_1}$.

With the vector $\mathbf{s}$ at hand, one can construct an estimate of the continuous field $f(x)$ given $\mathbf{s}$. How well can $f(x)$ be recovered based on $\mathbf{s}$? To make this question precise, we can find $\hat{f}(x|\mathbf{s})$: the minimum mean-square error (MMSE) estimate of $f(x)$ given $\mathbf{s}$. This is the estimate that will minimize the mean-square error between the original field and the reconstructed field given the observations $\mathbf{s}$. The error of this estimate will, of course, depend on the number, locations, and accuracies of the samples. For a given C_B, our objective is to choose the number of the samples M and the locations of the samples $\xi_1, ..., \xi_M$, while satisfying $C_T \leq C_B$, with the objective of minimizing the MMSE between $f(x)$ and $\hat{f}(x|\mathbf{s})$. We note that since the cost of each sample is assumed to be the same, by choosing the number of samples we also determine the cost of each sample.

This problem can be stated as one of minimizing over Δ_x, x_0, and M to determine the error:

$$\varepsilon(C_B) = \min_{\Delta_x, x_0, M} E\left[\int_D |f(x) - \hat{f}(x|\mathbf{s})|^2 dx\right], \tag{3}$$

subject to $C_T \leq C_B$. We consider all signals and estimators over some bounded domain D. Let $K_f(x_1, x_2) = E[f(x_1)f^*(x_2)]$ and $K_n(x_1, x_2) = E[n(x_1)n^*(x_2)]$ denote the covariance functions of $f(x)$ and $n(x)$, respectively. Here, $*$ denotes complex conjugation. We assume that $f(x)$ is a finite-energy random field, $\int_{-\infty}^{\infty} K_f(x,x)dx < \infty$, and $K_n(x,x)$ is bounded. An earlier version of this problem that does not include a system $\mathcal{L}$ or noise $n(x)$ was considered in [46].

At this point it is worth recalling some of the properties of MMSE estimation. As noted above, $\hat{f}(x|\mathbf{s})$ is the estimate that minimizes the mean-square error between $f(x)$ and $\hat{f}(x|\mathbf{s})$ for a given $\mathbf{s}$. The associated mean-square error $E[\int_D |f(x) - \hat{f}(x|\mathbf{s})|^2 dx]$ does not depend on the realization of the random vector $\mathbf{s}$, but only on the joint probability distribution of $f(x)$ and $\mathbf{s}$. Under the current problem formulation, for a given cost budget C_B, this joint probability distribution is determined by the number M and the locations of the samples $(\xi_1, ..., \xi_M)$. The formulation (3) seeks the best choices for these sampling strategy parameters.

The MMSE estimate in (3) can be written as $\hat{f}(x|\mathbf{s}) = \sum_{j=1}^M h_j(x) s_j = \mathbf{h}(x)\mathbf{s}$, where $\mathbf{h}(x) = [h_1(x), ..., h_M(x)]$ [49, Chap. 6]. We note that, given a set of samples, the set of functions $\mathbf{h}(x)$ are the optimal functions that minimize the mean-square error between the actual field and the reconstructed field. Here, $\mathbf{h}(x)$ satisfies the equation [49, Chap. 6]

$$\mathbf{K}_{f\mathbf{s}}(x) = \mathbf{h}(x)\mathbf{K_s}, \tag{4}$$

where $\mathbf{K}_{f\mathbf{s}}(x) = E[f(x)\mathbf{s}^\dagger] = [E[f(x)s_1^*], ..., E[f(x)s_M^*]]$ is the cross covariance between the input field $f(x)$ and the representation vector s, and $\mathbf{K_s} = E[\mathbf{s}\mathbf{s}^\dagger]$ is the autocovariance of $\mathbf{s}$. The symbol $\dagger$ denotes complex conjugate transpose. To determine the optimal linear estimate, one solves this last equation for $\mathbf{h}(x)$. The resulting estimate $\sum_{j=1}^M h_j(x) s_j$ can be interpreted as the orthogonal projection of the unknown random field $f(x)$ onto the subspace generated by the samples s_j, with $h_j(x)$ being the projection coefficients.

Before leaving this section, we would like to comment on the error introduced by estimating $f(x)$ only in the bounded region D. Let us make the following definitions: let $\hat{f}(x|\mathbf{s})$ be shortly denoted as $\hat{f}(x)$. Let us define $\hat{f}_D(x)$ as $\hat{f}_D(x) = \hat{f}(x)$ for $x \in D$ and $f_D(x) = 0$ for $x \notin D$. Then, the error of representing $f(x)$ with $\hat{f}_D(x)$ can be expressed as

$$E\left[\int_{-\infty}^{\infty} |f(x) - \hat{f}_D(x)|^2 dx \right]$$
$$= E\left[\int_{x \in D} |f(x) - \hat{f}_D(x)|^2 dx \right] + E\left[\int_{x \notin D} |f(x) - \hat{f}_D(x)|^2 dx \right] \tag{5}$$

$$= E\left[\int_{x \in D} |f(x) - \hat{f}_D(x)|^2 dx \right] + E\left[\int_{x \notin D} |f(x)|^2 dx \right] \tag{6}$$

$$= \varepsilon(C_B) + \int_{x \notin D} K_f(x, x) dx. \tag{7}$$

Hence, (7) states that the error of representing a field on the entire line can be expressed as the sum of two terms; the first term expressing the approximation error in this bounded region, and the second term expressing the error due to neglecting the function outside this bounded region (the energy of the field outside region D). Since the field is finite-energy, the second term can be made arbitrarily close to zero by taking a large enough region D and $\varepsilon(C_B)$ becomes a good measure of representation performance over the entire space.

3. RANDOM FIELD MODEL

In our examples, we use the GSM for $f(x)$. This is a random optical field model with various generalizations and applications [50–59]. A GSM source is characterized by the covariance function

$$K_f(x_1, x_2) = A_f \exp\left(-\frac{x_1^2 + x_2^2}{4\sigma_I^2} \right) \exp\left(-\frac{(x_1 - x_2)^2}{2\sigma_\nu^2} \right)$$
$$\times \exp\left(-\frac{jk}{2R}(x_1^2 - x_2^2) \right). \tag{8}$$

Here, $A_f > 0$, $j = \sqrt{-1}$. The parameters $\sigma_I > 0$ and $\sigma_\nu > 0$ determine the width of the intensity profile and the width of the complex degree of spatial coherence, respectively. R represents the wavefront curvature.

This covariance function may be represented in the form $K_f(x_1, x_2) = \sum_{k=0}^\infty \lambda_k \phi_k(x_1)\phi_k^*(x_2)$, where $\lambda_0 \geq \lambda_1 ... \lambda_k \geq \lambda_{k+1}, ...$ are the eigenvalues and $\phi_k(x)$ are the orthonormal eigenfunctions, $k \in Z_+$ [51,52,57]. This is the coherent-mode decomposition of the GSM source, where every eigenfunction is considered to correspond to one fully coherent mode. The ratio of the largest eigenvalue λ_n to the lowest eigenvalue λ_0 is given by $\lambda_n/\lambda_0 = (1/(\beta^2 + 1 + \beta[(\beta/2)^2 + 1]^{0.5}))^n$, where β is defined as

$$\beta = \frac{\sigma_\nu}{\sigma_I}. \tag{9}$$

β may be considered as a measure of the degree of (global) coherence of the field [52,57]. As β increases, the eigenvalues decay faster, so that the effective number of modes required to represent the field decreases and the field is said to be more coherent. In contrast, as β decreases, the eigenvalues decay slower, so that the effective number of modes required to represent the field increases and the field is said to be more incoherent.

Various aspects of the propagation of GSM fields through optical systems have been well studied; see, for instance [53–57,59]. A fundamental result in this area that we will make use of is the following: say we have an optical system that may be represented by an $ABCD$ matrix (ray-transfer matrix). When a GSM field passes through such an optical system, the output is again a GSM field with new parameters σ_I', σ_ν', and R_{out}' [53,54]. It is known that the ratio $\beta = \sigma_I'/\sigma_\nu'$ does not change as the field passes through such systems [50,53,54]. Hence, σ_ν' is given simply by $\sigma_\nu' = \beta\sigma_I'$.

Our system noise model is characterized by the following covariance function: $K_n(x_1 - x_2) = A_n \exp(-(x_1 - x_2)^2/2\sigma_{\nu,n}^2)$, with $\sigma_{\nu,n} = \beta_n\sigma_I$, $\beta_n < \beta$.

4. TRADE-OFF CURVES FOR GSM FIELDS ARE INVARIANT UNDER PROPAGATION THROUGH FIRST-ORDER OPTICAL SYSTEMS

We now consider the problem of sampling the output of a first-order optical system in order to represent the input optical field. Such systems are also referred to as $ABCD$ systems or quadratic-phase systems [60]. Mathematically represented by linear canonical transforms [61], these systems encompass arbitrary concatenations of lenses, mirrors and sections of free space, as well as quadratic graded-index media. In the next section, we will consider a given bit budget and find the minimum possible representation error for that budget. Varying the bit budget, we will obtain trade-off curves
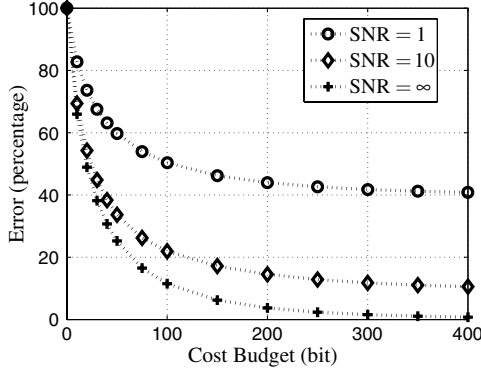
Fig. 1.   Error versus cost budget C_B, $\beta = 1/8$ (varying SNR).

between the error and the cost budget (for instance, look forward to Fig. 1 for an example). Here, we are concerned with how first-order optical systems change these trade-off curves; in other words, does it make any difference if we represent the signal with samples of the output of such a system, rather than with samples of the input itself? We will show that for GSM fields, the cost-error curves are invariant under passage through arbitrary *ABCD* systems; that is, these systems have no effect on the error versus cost trade-off curves. Moreover, we show that the optimum sampling strategy at the output can be easily found by scaling the optimum sampling strategy at the input. We assume that the parameters A, B, C, D of the *ABCD* matrix are real with $AD - BC = 1$. We first consider the case where there is no system noise $n(x)$, and then discuss the effects of noise.

Let us express the covariance function associated with a GSM field with parameters σ_I, β, R as

$$K_{\sigma_I,\beta,R}(x_1,x_2) = A_f \exp\left(-\frac{x_1^2 + x_2^2}{4\sigma_I^2}\right)\exp\left(-\frac{(x_1 - x_2)^2}{2(\beta\sigma_I)^2}\right)$$
$$\times \exp\left(-\frac{jk}{2R}(x_1^2 - x_2^2)\right). \tag{10}$$

We note the following scaling property for the $R = \infty$ case:

$$K_{\sigma_I',\beta,\infty}(-x_1,x_2) = K_{\sigma_I,\beta,\infty}\left(-x_1\frac{\sigma_I}{\sigma_I'}, x_2\frac{\sigma_I}{\sigma_I'}\right), \tag{11}$$

which expresses the fact that the covariance function associated with a given σ_I can be found by scaling that associated with another σ_I'. The error expression depends on the joint distribution of the samples **s** and the field $f(x)$, which in turn is determined through the covariance functions. Considering the representation of $f(x)$ in terms of its samples, we also note that for a given set of σ_{m_i}, the cost associated with a set of sampling points remains unchanged if the sampling points are scaled by σ_I'/σ_I. Hence, we conclude that for the case $R = \infty$ and $\mathcal{L}$ is the identity, the error does not depend on σ_I, provided the sampling points are scaled appropriately. As a result, the cost-error trade-off curves will be the same for different values of σ_I, and the optimum sampling strategies will be scaled versions of each other.

Here we show that the conclusion of the preceding paragraph continues to remain valid even when $R \neq \infty$. We will first show that for a given set of sampling points $\xi_1, ..., \xi_M$, and a given covariance matrix K_m, the associated costs

and the error for all values of R are the same. This, in fact, stems from the fact that the curvature term corresponds to uncorrelated phase terms. Let the covariance function associated with $f(x)$ be $K_{\sigma_I,\beta,\infty}(x_1,x_2)$. Let $\bar{f}(x)$ be the zero-mean complex proper field with the covariance function

$$E[\bar{f}(x)\bar{f}^*(x)] = K_{\sigma_I,\beta,R}(x_1,x_2) \tag{12}$$

$$= K_{\sigma_I,\beta,\infty}(x_1,x_2)\exp\left(-\frac{jk}{2R}(x_1^2 - x_2^2)\right) \tag{13}$$

$$= E[f(x)f^*(x)]\exp\left(-\frac{jk}{2R}(x_1^2 - x_2^2)\right). \tag{14}$$

We first observe that the presence of a curvature does not affect the cost associated with a sample. The cost associated with the ith sample $\bar{s}_i = \bar{f}(\xi_i) + m_i$ is given by $C_{\bar{s}_i} = \log_2(\sigma_{\bar{s}_i}^2/\sigma_{m_i}^2)$, where $\sigma_{\bar{s}_i}^2 = E[|\bar{s}_i|^2] = E[|\bar{f}(\xi_i)|^2] + E[|m_i|^2] = E[|f(\xi_i)|^2] + E[|m_i|^2]$. Hence, the cost of a sample with a given $E[|m_i|^2] = \sigma_{m_i}^2$ does not depend on the value of R.

We now show that the error does not depend on the value of R; that is, for a given set of sampling locations and a given set of σ_{m_i}, the errors associated with estimating $f(x)$ and $\bar{f}(x)$ are the same. Let us define the vector **g** as $\mathbf{g} = [f(\xi_1), ..., f(\xi_M)]^T$, $i = 1, ..., M$. Now, the vector of finite accuracy samples of $f(x)$ is given by $\mathbf{s} = \mathbf{g} + \mathbf{m}$, where $\mathbf{m} = [m_1, ..., m_M]^T$. Let the $M \times M$ covariance matrix of the finite accuracy samples be denoted by $E[\mathbf{ss}^\dagger] = K_\mathbf{s} = K_\mathbf{g} + K_\mathbf{m}$, where the element in the ith row and lth column of $K_\mathbf{g}$ is given by $K_{\sigma_I,\beta,\infty}(\xi_i,\xi_l)$, $i,l = 1, ..., M$. The cross covariance between $f(x)$ and s is given by the $1 \times M$ row vector $E[f(x)\mathbf{s}^\dagger] = \mathbf{d}(x)$, where the lth element is given by $K_{\sigma_I,\beta,\infty}(x,\xi_l)$. Similarly, we define $\bar{\mathbf{s}} = \bar{\mathbf{g}} + \bar{\mathbf{m}}$, where $\bar{\mathbf{g}} = [\bar{f}(\xi_1), ..., \bar{f}(\xi_M)]^T$. Consequently, we have $K_{\bar{\mathbf{s}}} = K_{\bar{\mathbf{g}}} + K_{\bar{\mathbf{m}}}$, where the element in the ith row and lth column is given by $K_{\sigma_I,\beta,R}(\xi_i,\xi_l)$, and $E[\bar{f}(x)\bar{\mathbf{s}}^\dagger] = \bar{\mathbf{d}}(x)$, where the lth element is given by $K_{\sigma_I,\beta,R}(x,\xi_l)$. Now, let $T = \mathrm{diag}(t_i)$, $t_i = \exp(-(jk/2R)\xi_i^2)$, $i = 1, ..., M$. We observe that

$$K_{\bar{\mathbf{s}}} = K_{\bar{\mathbf{g}}} + K_\mathbf{m} \tag{15}$$

$$= TK_\mathbf{g}T^\dagger + TK_\mathbf{m}T^\dagger \tag{16}$$

$$= TK_\mathbf{s}T^\dagger, \tag{17}$$

where (16) follows from the fact that $TK_\mathbf{m}T^\dagger = \mathrm{diag}(t_i)\mathrm{diag}(\sigma_{m_i}^2)\mathrm{diag}(t_i^*) = \mathrm{diag}(\sigma_{m_i}^2) = K_\mathbf{m}$, since $|t_i| = |\exp(-(jk/2R)\xi_i^2))| = 1$. We also observe that

$$\bar{\mathbf{d}}(x) = \exp(-(jk/2R)x^2))\mathbf{d}(x)T^\dagger. \tag{18}$$

Now, using these results, we finally show that the error is independent of the value of R. We consider the error for the field at a given point x. Denoting the MMSE estimate of $\bar{f}(x)$ given $\bar{\mathbf{s}}$ as $\hat{\bar{f}}(x|\bar{\mathbf{s}})$, the associated MMSE can be expressed as [49, Chap. 6]

$$E[|\bar{f}(x) - \hat{\bar{f}}(x|\bar{\mathbf{s}})|^2] = K_{\sigma_I,\beta,R}(x,x) - \bar{\mathbf{d}}(x)K_{\bar{\mathbf{s}}\bar{\mathbf{s}}}^{-1}\bar{\mathbf{d}}(x)^\dagger \tag{19}$$

$$= K_{\sigma_I,\beta,\infty}(x,x) - \mathbf{d}(x)K_{\mathbf{ss}}^{-1}\mathbf{d}(x)^\dagger \tag{20}$$

$$= E[|f(x) - \hat{f}(x|\mathbf{s})|^2]. \tag{21}$$

In obtaining (20), we used (17), (18), $TT^\dagger = I$, and $|\exp(-(jk/2R)\xi_i^2))| = 1$, where I is the $M \times M$ identity matrix. Hence, we have shown that the value of R does not change the error.

So far we have shown that (i) for $R = \infty$, the error does not depend on σ_I, provided the sampling points are appropriately scaled; (ii) for a given set of sampling points and σ_{m_i}, the associated errors and costs do not depend on R. Thus, we conclude that for a given GSM field with a specified value of β, the cost-error trade-off curves associated with the problem of estimating a field based on its own samples do not depend on σ_I and R. Now, recall that GSM fields remain GSM fields with the same β, but different σ_I and R after passing through first-order optical systems [50,53,54]. This, combined with the previous observations, show that the error associated with estimating the output field by sampling the output, is the same as the error associated with estimating the input field by sampling the input (under the same cost).

Finally, we consider the problem of sampling the output of a first-order optical system in order to estimate the input field. We first recall that the MMSE is invariant under unitary transformations; that is, the MMSE associated with estimating $f(x)$ based on a random vector s is the same as the MMSE associated with estimating $\mathcal{L}\{f(x)\}$, if $\mathcal{L}$ is a unitary transformation. We also recall that optical systems represented by real A, B, C, D parameters are unitary systems [62, Chap. 9]. Hence, for any such system, the MMSE associated with estimating the input of the optical system and the output of the optical system based on a given set of samples of the output are the same. Thus, combining this with the observations of the previous paragraph, we conclude that the error versus cost trade-offs for the estimation of the input of an optical system based on the samples of the input field are the same as those based on the samples of the output field. (The same conclusion also holds for estimating the output based on the samples of the input or the output.) In other words, finite-accuracy samples of the output field are as good as finite-accuracy samples of the input field for the broad class of first-order optical systems.

We now discuss the effect of noise $n(x)$. Here we will show that, as in the noiseless case, when the system $\mathcal{L}$ is identity and $R = \infty$, the error value does not depend on σ_I, provided the sampling points are scaled appropriately. To show this in the noisy case, we need to show that the associated covariance functions can be scaled with σ_I. (i) The scaling property of $K_{f(x)}$ was already discussed at the beginning of this section. (ii) The noise covariance function also scales with σ_I, in a manner similar to (11). It follows from (i) and (ii) that the covariance of the observations also scales with σ_I. We also note that, due to statistical independence of $f(x)$ and $n(x)$, the cross covariance of $f(x)$ and s only depends on the covariance function of $f(x)$, which is known to scale with σ_I. Hence, all associated covariances have the scaling property. Thus, we can now conclude that the error for a given set of sampling points for a given σ_I, can be found by looking at the error for another σ_I at a scaled set of sampling points. We also note that for a given set of σ_{m_i}, the cost associated with a set of sampling points, remains unchanged under appropriate scaling. This implies that the trade-off curves are invariant for different σ_I values and the optimum sampling points can be found by scaling.

5. OPTIMAL TRADE-OFFS BETWEEN REPRESENTATION ERROR AND BIT BUDGET

In this section, we present trade-off curves between the error and the cost budget, and the optimum sampling parameters achieving these curves.

To be able to observe the effect of the degree of coherence on the results, we use two different β values: $\beta = 1/8$ and $\beta = 1$ in our numerical experiments. We choose $\beta_n = 1/32$. We consider different noise levels parameterized through the SNR, defined as the ratio of the peak signal and noise levels: $\text{SNR} = A_f/A_n$. We consider the values $\text{SNR} = 1, 10, \infty$ to cover a wide range of problem instances. We choose the interval D equal to $[x_L, x_H] = [-5\sigma_I, +5\sigma_I]$ to ensure that the signal values are safely negligible outside D. We report the error as a percentage defined as $100\varepsilon(C_B)/\varepsilon_0$, where $\varepsilon_0 = \int_{-\infty}^{\infty} K_f(x,x)\mathrm{d}x = A_f\sqrt{2\pi}$.

We would like to note that error-cost trade-off curves do not depend on the total energy of the signal. More precisely, when there is no system noise $n(x)$, the error-cost curves are independent of the constant A_f in (10). When there is system noise $n(x)$, the error-cost curves do not depend on the individual values of A_f and A_n, but only on the ratio $\text{SNR} = A_f/A_n$.

Based on the discussion of Section 4, we note that in the noiseless case ($\text{SNR} = \infty$), the presented cost-error trade-off curves are valid for any $ABCD$ system with real parameters, $AD - BC = 1$. The optimum sampling points are easily found by scaling in proportion to the ratio of input and output σ_I s. When $\text{SNR} \neq \infty$, the curves are obtained for the case $\mathcal{L}$ is the identity operator and $R = \infty$, and these do not generalize to arbitrary $ABCD$ systems. But the optimum sampling points for one value of σ_I can still be found from those for another by scaling.

For simplicity in presentation, in our simulations we focus on Δ_x and set the less interesting $x_0 = 0$. To compute the error expressions and optimize over the parameters of the representation strategy, we discretize the x space with the spacing Δ_c. We approximate the integral in (3) as $\sum_{k \in D_N} |f(k\Delta_c) - \hat{f}(k\Delta_c|s)|^2\Delta_c$, where $D_N = \{k : k\Delta_c \in D\}$. The estimates are only calculated at these discrete points: $\hat{f}(k\Delta_c|s) = \mathbf{h}(k\Delta_c)s$. To determine the estimate functions $\mathbf{h}(k\Delta_c)$, we solve the equation $\mathbf{K_{fs}}(k\Delta_c) = \mathbf{h}(k\Delta_c)\mathbf{K_s}$ for each $k \in D_N$. In order to find the optimum sampling interval, we use a brute force method, where for a given C_B we calculate the error for varying Δ_x and M, and choose the values providing the least error. We note that the optimization variable Δ_x and the discretization variable Δ_c are not the same. Δ_x is the sampling interval whose optimal value we seek, whereas Δ_c is the discrete grid spacing we employ in the numerical experiments.

Figures 1 and 2 present the error versus bit budget curves for varying SNR for a relatively incoherent field ($\beta = 1/8$) and for a relatively coherent field ($\beta = 1$), respectively. As expected, the error decreases with increasing cost budget in all cases. We note that $\varepsilon(C_B)$ is very sensitive to increases in C_B for smaller C_B. Then it becomes less responsive and eventually saturates.

We observe that in each of these figures, as the noise level becomes higher, it becomes more difficult to obtain low values of error. We observe that for both values of β, when there is no system noise ($\text{SNR} = \infty$), the error goes to zero as we increase the cost. This means that, no matter how small the

error tolerance $\varepsilon > 0$ is specified to be, the continuous finite-energy field can be represented with a finite number of bits. This observation is discussed in more detail in Section 7.

Comparing these figures, we observe that for the relatively incoherent case (Fig. 1), it is more difficult to achieve low values of error for a given bit budget. But as the field becomes more coherent (Fig. 2), the field values at different locations become more correlated with each other, the total uncertainty in the field decreases, and it becomes a lot easier to achieve lower values of error.

We now investigate the relationship between the optimum sampling strategies and the problem parameters C_B, SNR, and β. The optimum sampling interval Δ_x and the optimum number of samples M that achieve the errors given in Fig. 1 are presented in Figs. 3 and 4 for SNR $= \infty$ and SNR $= 1$. The optimum sampling interval Δ_x and the optimum number of samples M that achieve the errors given in Fig. 2 are presented in Figs. 5 and 6 for SNR $= \infty$ and SNR $= 1$.

When there is no system noise $n(x)$, the optimum sampling strategies can be informally interpreted in the light of the competition between the following driving forces: (i) to have as many effectively uncorrelated samples as possible, (ii) to have samples whose variances are as high as possible, and (iii) to have samples which are as highly accurate as possible. When there is system noise $n(x)$, each sample tells less about the value of the field. In order to wash out the effect of noise, one is often willing to take samples at field locations which are considerably correlated, and which one would probably not take samples at, had there been no noise.

We observe that in all cases, in general, as C_B increases, the optimum sampling interval decreases and the number

of samples increases: when we have more bits to spend, we use a higher number of more closely spaced samples. When C_B is low, the optimal strategy is to use a low number of more distantly spaced samples so that each sample has a reasonable accuracy and each of them provides effectively new information about the field. As the allowed cost increases, we can afford more samples with high enough accuracies and we prefer to use more closely spaced samples so that we can get more information about field values we previously had to neglect when the allowed cost was lower.

Comparing Figs. 3 and 4 (or Figs. 5 and 6), we observe that as the noise level increases, the samples should be taken more closer (the sampling interval decreases). When a sample is noisy, one would expect the information provided by that sample to be smaller, encouraging us to take more closely spaced samples so as to compensate for the effects of noise. We also observe that as the noise level increases, one should take a higher number of samples M. This observation may seem trivial, since decreasing the sampling interval automatically increases the number of samples within a certain spatial range. However, we note that here the range over which samples are taken does not remain constant but also decreases. (The variances of field values decrease as we move away from the $x = 0$ point, so that the field here is highly contaminated by noise. Since samples taken here are of little value for representing the field, it is reasonable to expect that it will be better not to take these samples, thereby decreasing the spatial range the samples are taken over.) However, the decrease in the spatial range is not as much as to compensate the decrease in the sampling interval, so in the end the number of samples taken increases.


Fig. 2.   Error versus cost budget C_B, $\beta = 1$ (varying SNR).


Fig. 4.   Number of samples and optimum sampling interval versus cost budget, $\beta = 1/8$, SNR $= 1$.
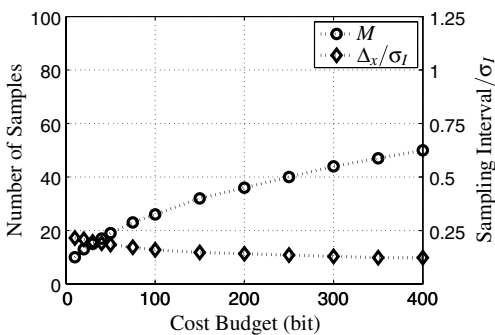

Fig. 3.   Number of samples and optimum sampling interval versus cost budget, $\beta = 1/8$, SNR $= \infty$.

Fig. 5.   Number of samples and optimum sampling interval versus cost budget, $\beta = 1$, SNR $= \infty$.

Fig. 6.   Number of samples and optimum sampling interval versus cost budget, $\beta = 1$, SNR = 1.

Comparing Figs. 3 and 5, we see that when the field is more coherent, it is desirable to take a fewer number of samples which are farther apart. When the field is more coherent, under the GSM correlation structure, the field value at each point becomes more correlated with field values farther away. Hence, there is a tendency to space the samples well in order to get effectively new information from each sample. Also, the variances of the field values decrease as we move further away from the $x = 0$ point, so we prefer not to waste any of our bit budget on such samples. As a result, the optimum number of samples is smaller, which is consistent with the fact that more coherent fields have a lower number of effective modes (the number of uncorrelated random variables required to effectively represent the field).

6. COMPARISON WITH SHANNON–NYQUIST SAMPLING BASED APPROACHES

A common approach in sampling optical fields is to use the Shannon–Nyquist sampling theorem as a guideline. As outlined in Section 1, in this traditional approach, one determines an effective frequency extent B, and an effective spatial extent L, which are used to determine the sampling interval, and the spatial extent the samples will be taken over, respectively. Here, we will compare the error versus cost budget curve that is obtained following this traditional approach with the optimal curves obtained with our approach and shown in Figs. 1 and 2. But first we review how the traditional Shannon–Nyquist approach applies to random fields. A fundamental result in this area states that the Shannon–Nyquist sampling theorem can be generalized to wide-sense stationary (WSS) signals: a band-limited WSS signal can be reconstructed in the mean-square sense from its equally spaced samples taken at the Nyquist rate [63]. References [64,65] further generalize this result to nonstationary random fields: let $v(x) \in \mathbb{R}$ be a finite-energy random field. Let us consider the covariance function of the Fourier transform of the field defined as $S_v(\nu_1, \nu_2) = E[V(\nu_1)V^*(\nu_2)]$, where $V(\nu)$ is the Fourier transform of $v(x)$. If $S_v(\nu, \nu) = 0$, for $|\nu| > B/2$, then the field can be recovered from its samples in the mean-square sense; that is, $E[|v(x) - \sum_{k=-\infty}^{\infty} v(k/B)\text{sinc}(xB - k)|^2] = 0$.

We now explicitly work out the conventional sampling approach for GSM fields. The effective spatial extent of the field will be determined by looking at the intensity distribution $K_f(x, x) = \exp(-x^2/2\sigma_I^2)$, which has a Gaussian profile with standard deviation σ_I. Most of the energy of a Gaussian lies within a few standard deviations so that the effective spatial

extent can be taken as $[-r\sigma_I, r\sigma_I]$. The intensity of the Fourier transform of the field; that is, the diagonal of the covariance function of the Fourier transform of the field also has a Gaussian profile $S_f(\nu, \nu) \propto \exp(-f^2/2\sigma_{I,F}^2)$, where $S_f(\nu_1, \nu_2) = E[F(\nu_1)F^*(\nu_2)]$, where $F(\nu)$ is the Fourier transform of $f(x)$, and $\sigma_{I,F} = (1/2\pi)\sqrt{(1/\beta^2) + (1/4)}/\sigma_I$ (see, for instance [66]). We take the effective frequency extent as $[-r\sigma_{I,F}, r\sigma_{I,F}]$. This implies a sampling interval of $1/(2r\sigma_{I,F})$. The number of samples is found by dividing the effective spatial extent to the sampling interval as follows:

$$M_s = \frac{2r\sigma_I}{1/(2r\sigma_{I,F})} = \frac{2r^2}{\pi}\left(\frac{1}{\beta^2} + \frac{1}{4}\right)^{0.5}. \qquad (22)$$

Hence, for each cost budget value C_B, the cost associated with each sample will be C_B/M_s. In our experiments, we round M_s to the nearest integer. To ensure a fair comparison with the approach of this paper, we again use the mean-square estimate to estimate the signal from the Nyquist samples.

We now compare the error versus bit budget trade-offs obtained with the approach presented in this article, with those obtained by using the traditional approach described above. We use two different r values; $r = 2$, and $r = 3$. Figures 7 and 8 compare the trade-off curves for $\beta = 1/8$ and $\beta = 1$, respectively. The dotted curves and the dashed/solid lines show the results for the optimal sampling scheme and the sampling theorem based scheme, respectively. As expected, for all cases, the optimum sampling strategy gives better trade-offs compared to the sampling strategies based on the sampling theorem.

We note that when there is no system noise $n(x)$, and if we determine the effective extents appropriately, we would expect to obtain error values close to zero for high values of cost budget. We observe that this is indeed the case for $r = 3$, but not for $r = 2$. This suggests that $r = 2$ is a poor choice for defining the effective extents, and illustrates the importance of determining effective extents appropriately.

When $r = 3$ and there is no system noise, for both relatively low and high degrees of coherence, the optimal strategy and the traditional strategy differ by a greater amount for low values of cost budget. This observation may be interpreted as follows: when the cost budget is low, the relatively high number of samples dictated by the sampling theorem will result in the samples being relatively inaccurate, leading to poor performance. (As we have seen earlier, for low cost
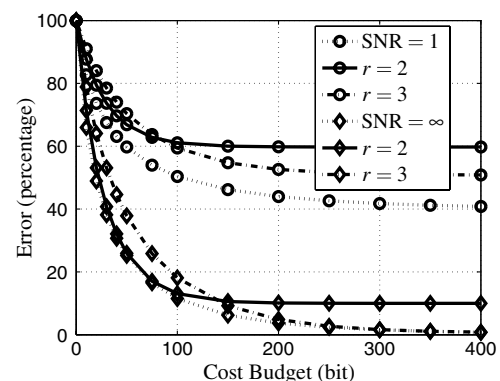

Fig. 7.   Error versus cost budget C_B, $\beta = 1/8$ (varying SNR). The dotted lines are for optimal sampling strategies and the corresponding dashed and solid lines are for sampling theorem-based strategies.

Fig. 8. Error versus cost budget C_B, $\beta = 1$ (varying SNR). The dotted lines are for optimal sampling strategies and the corresponding dashed and solid lines are for sampling theorem-based strategies.
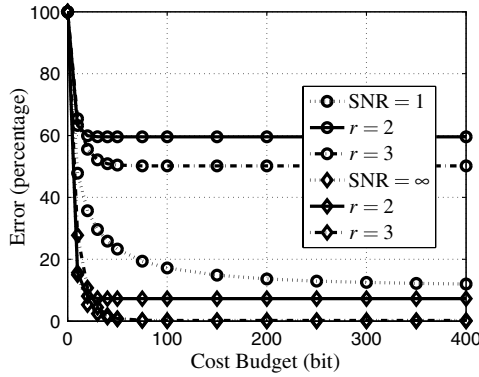
values, it is better to use a smaller number of samples with relatively better accuracy.) As the cost budget increases, the difference between the two approaches gets smaller, and both strategies achieve error values very close to 0, as expected. For low values of cost budget, the traditional approach with $r = 2$ dictates a sampling strategy closer to the optimal one, compared to $r = 3$, and gives error values closer to the optimal strategy. Yet, as observed above, it gives relatively poor error values for higher values of cost budget, and therefore cannot be considered a good sampling approach for all values of the cost budget.

When the system noise level is high, the difference between the optimal and traditional strategies is pronounced for almost all values of cost budget. The sampling theorem assumes that the samples will be noiseless, and therefore cannot exploit the opportunity for noise elimination through oversampling. (We observe that the traditional strategy with $r = 2$ gives poorer results compared to the $r = 3$ case, which may be attributed to the relatively low number of samples dictated by the former.) We observe that the performance difference between the traditional approaches and the optimal strategy is more pronounced for the coherent case. When the field is more coherent, the sampling theorem-based strategy dictates the use of a fewer number of more distantly spaced samples, compared to the incoherent case. However, in the presence of noise, the optimal strategy is not that much different for the incoherent and coherent cases, and dictates that we use a comparably larger number of more closely spaced samples even when the field is coherent. Therefore, the traditional sampling strategies are more markedly inferior than the optimum strategy in the coherent case.

7. DISCUSSION

In Section 5 we illustrated that, given an arbitrarily small but nonzero error tolerance, it is possible to represent a finite-energy optical field with a finite number of bits without exceeding that error tolerance. At first glance, this may appear as a surprising observation. After all, the optical field in question takes continuous amplitude values in continuous and unbounded space, and attempting to use a finite numbers of bits to represent such a field is a severe restriction: such finite representations usually involve a finite number of samples each quantized to a finite number of levels. Therefore, here we further discuss this from different perspectives.

First, consider the very simple case of a single sample of the field. Let us assume this sample can assume values between A_{low} and A_{high} and we have agreed to represent this value with an error tolerance of ΔA. Then, it follows that there will be $\sim (A_{\text{high}} - A_{\text{low}})/\Delta A$ distinguishable levels, which can then be represented by $\sim \log_2((A_{\text{high}} - A_{\text{low}})/\Delta A)$ bits.

Now let us return to the field $f(x)$. We may think of the finite-energy condition as a limitation on how large the amplitude values of the field can be. On the other hand, the specified error tolerance can be considered to determine the minimum separation of two signals such that they are still considered distinguishable. The finite-energy condition restricts the signal to reside within a hypersphere of specific radius, whereas the error tolerance defines a certain volume within which signals are considered indistinguishable. Roughly speaking, the number of distinguishable signals is given by the volume of the hypersphere divided by the volume defined by the finite error tolerance. Since this number is finite, the signal can be represented by a finite number of bits.

We now take a somewhat more mathematical, closer look at this issue. In each step of our argument, we introduce a limitation, a form of "finiteness," in the representation of the field (such as limiting the fields to a bounded region), and argue that the error introduced by each of these limitations can be made arbitrarily small. This way, we aim to illustrate how different forms of "finiteness" contribute to the overall picture. Our approach is based on the coherent-mode decomposition. We also note that it is more common to discuss concepts related to "finiteness" in a deterministic setting, and in connection with band-limited approximations, rather than the stochastic setting and approximations based on covariance functions we employ.

Let us consider a finite-energy zero-mean random field that will be approximated using a finite number of bits. For the sake of convenience, let us assume that the random field takes real values. Let us first focus on the error introduced by the limitation of representing the signal in a bounded region D instead of the infinite line. As stated in (6), the total error of such an approximation can be expressed as the sum of two terms: one is the approximation error on D, and the other one is the energy outside D. The energy outside D can be made arbitrarily small by taking D large enough. This is the first form of "finiteness" introduced in the representation of the signal.

We now focus on the approximation error on D. The question is whether it is possible to make the approximation error arbitrarily close to zero; that is, whether it is possible to represent the field in a bounded region with a finite number of bits. The answer is not obvious since we are dealing with a field taking continuous amplitude values on a bounded but still continuous space. To give an affirmative answer, we will rely on the existence of the Karhunen–Lóeve expansion of the covariance function of the unknown field with a discrete eigenvalue spectrum as

$$K_f(x_1, x_2) = \sum_{k=0}^{\infty} \lambda_k \phi_k(x_1) \phi_k^*(x_2), \qquad (23)$$

where $\lambda_0 \geq \lambda_1 \ldots \lambda_k \geq \lambda_{k+1}, \ldots$ are the eigenvalues and $\phi_k(x)$ are the orthonormal eigenfunctions, $k \in Z_+$. This is the coherent-mode decomposition of the random optical field. Here, each λ_i and ϕ_i pair is considered to correspond to one fully coherent

mode. Existence of such an expansion for covariance functions on a bounded region is guaranteed by Mercer's theorem; see, for example, [67, Chap. 7]. Therefore, the signals can be decomposed as

$$f(x) = \sum_{k=1}^{\infty} z_k \phi_k(x), \qquad x \in D, \qquad (24)$$

where the random variables z_k are zero-mean random variables with $E[|z_k|^2] = \lambda_k$. Hence, a continuous field on the bounded region can be represented with an infinite but at least denumerable number of variables, namely the random variables z_k, $k \in Z_+$. Here, it is also known that $\int_D K_f(x_1, x_2) dx = \sum_{k=0}^{\infty} \lambda_k$ [67, Chap. 7]. Since $K_f(x_1, x_2)$ is finite-energy, the left-hand side of this equation (the energy on the region D), is also finite. Hence, the right-hand side is also finite and we should have $\lambda_k \to 0$ as $k \to \infty$. Now, let us consider the truncation error

$$E\left[\int_D \left| f(x) - \sum_{k=1}^{N} z_k \phi_k(x) \right|^2 dx \right]$$

$$= E\left[\int_D \left| \sum_{k=1}^{\infty} z_k \phi_k(x) - \sum_{k=1}^{N} z_k \phi_k(x) \right|^2 dx \right] \qquad (25)$$

$$= E\left[\int_D \left| \sum_{k=N+1}^{\infty} z_k \phi_k(x) \right|^2 dx \right] \qquad (26)$$

$$= \sum_{k=N+1}^{\infty} E[|z_k|^2] \qquad (27)$$

$$= \sum_{k=N+1}^{\infty} \lambda_k. \qquad (28)$$

Thus, by choosing larger and larger but still finite values of N, we can bring the truncation error below any finite value, no matter how small. This observation shows that finite-energy optical fields can be represented by a finite number of variables $(z_1, ..., z_N)$ for any given nonzero error tolerance.

Finally, we would like to argue that it is possible to represent the field not only with a finite number of variables, but also with a finite number of bits. Here the question is whether it is possible to represent the finite-variance random variables $z_1, ..., z_N$ with a finite number of bits, to meet a given arbitrarily small nonzero error tolerance. The answer is affirmative and a classical result in information theory (rate-distortion theory [68, Chap. 13]). Although one would need an infinite number of bits to represent a continuous number perfectly (with zero error), it is possible to represent such a number with a finite number of bits with an arbitrarily small but nonzero error. With this last step, we conclude our argument showing that finite-energy optical fields can be represented by a finite number of bits with an arbitrarily small nonzero error tolerance.

In the first step of the argument of this section, we argued that the error introduced by limiting the signal to a bounded region can be made small. Actually, this step can be dispensed with altogether since finite-energy fields have Karhunen–Loéve expansions on the infinite line with a discrete eigenvalue spectrum (and hence coherent-mode decompositions with denumerable modes). Indeed, in the literature authors sometimes write the coherent-mode decomposition of an optical field in the form of a summation without explicit reference to a bounded domain or any detailed discussion of the existence of such an expansion on the infinite line. Here we would like to point out that this practice is supported by mathematical results: [69, Thm. 1] states that along with continuity, having $\int_{-\infty}^{\infty} K_f(x, x) dx < \infty$ and $K_f(x, x) \to 0$ as $|x| \to \infty$ is sufficient to ensure such a representation. We note that both of these conditions are plausible in a physical context: the first one is equivalent to the finite-energy assumption and the second one requires the intensity of the field to vanish as $|x|$ increases, properties one commonly expects from physically realizable fields.

8. CONCLUSIONS

Although optical fields are usually represented by functions of continuous variables, we know that in effect they carry a finite amount of information. This finiteness is intrinsically related to the finiteness of the energy and the specified nonzero error tolerance or noise in the system. Since we would like to quantify these, and since we often use digital systems to process information, we have set ourselves the goal of representing the field as efficiently as possible; that is, with as small a number of bits as possible.

We focused on various trade-offs in the representation of optical fields, mainly: (i) the trade-offs between the achievable error and the cost budget, (ii) the trade-offs between the accuracy, spacing, and number of samples. We have derived the optimal bounds for simultaneously achievable bit cost and error and obtained the optimal sampling parameters necessary to achieve them. These performance bounds are not only of interest for better understanding of information relationships inherent in propagating wave fields, but can also lead to guidelines in practical scenarios.

In contrast to common practice, which often treats sampling and quantization separately, we have explicitly focused on the interplay between limited spatial resolution and limited amplitude accuracy. Under a given cost budget, we have investigated whether it is better to take a higher number of samples with relatively lower cost per sample (hence, with lower amplitude accuracy), or a lower number of samples with relatively higher cost per sample (hence, with higher amplitude accuracy). Our results further reveal how, for the optimum number of samples determined, we should choose the space and frequency coverages. That is, we have answered the question of whether it is better to take more closely spaced samples (with wider frequency coverage but smaller spatial coverage), or to take more distant samples (with smaller frequency coverage but larger spatial coverage). We have seen that in certain cases, sampling at rates different than the Nyquist rate turns out to be more efficient. We also investigated how these results are affected by the degree of coherence of the field and the noise level. Furthermore, we observed how the optimal sampling parameters change with increasing cost budget.

The optical field at one part of a system is not independent from the optical field at another part of the system. In other words, knowledge of the field at one part of the system gives

us a certain degree of information about the field at other parts. Thus, we also considered the case where the signal is represented by samples taken after the signal passes through a linear system. For the case of GSM fields, when there is no noise, we have shown that finite-accuracy samples of the output field are as good as finite-accuracy samples of the input field, for the broad class of first-order optical systems. This class includes arbitrary concatenations of lenses, mirrors and sections of free space, as well as quadratic graded-index media. The cost-error trade-off curves obtained turn out to be the same as those obtained for direct sampling of the input, and the optimum sampling points can be found by a simple scaling of the direct sampling results.

Our results may be of use in practical applications where partially coherent information-bearing light fields are to be recorded. Our conclusions may be particularly relevant to digital holography where measurement devices located at discrete points in space are used to record the hologram. Our results suggest that in some cases it may be more efficient to sample the optical field at rates different than the Nyquist rate. Our bit budget framework can also be of interest in a typical holography scenario, where one first records a huge amount of data, and then proceeds to process this data in order to transmit or store it using a smaller number of bits. With our approach, instead of this two-step procedure, one can acquire the data more efficiently in the first place, under the given bit budget.

ACKNOWLEDGMENTS

REFERENCES

1. G. T. D. Francia, "Resolving power and information," J. Opt. Soc. Am. **45**, 497–501 (1955).
2. D. Gabor, "Light and information," in *Progress in Optics*, E. Wolf, ed. (Elsevier, 1961), Vol. **I**, pp. 109–153.
3. W. Lukozs, "Optical systems with resolving powers exceeding the classical limit," J. Opt. Soc. Am. **56**, 1463–1472 (1966).
4. W. Lukozs, "Optical systems with resolving powers exceeding the classical limit II," J. Opt. Soc. Am. **57**, 932–941 (1967).
5. G. T. Di Francia, "Degrees of freedom of an image," J. Opt. Soc. Am. **59**, 799–804 (1969).
6. F. Gori and G. Guattari, "Effects of coherence on the degrees of freedom of an image," J. Opt. Soc. Am. **61**, 36–39 (1971).
7. F. Gori and G. Guattari, "Shannon number and degrees of freedom of an image," Opt. Commun. **7**, 163–165 (1973).
8. A. Starikov, "Effective number of degrees of freedom of partially coherent sources," J. Opt. Soc. Am. **72**, 1538–1544 (1982).
9. G. Newsam and R. Barakat, "Essential dimension as a well-defined number of degrees of freedom of finite-convolution operators appearing in optics," J. Opt. Soc. Am. A **2**, 2040–2045 (1985).
10. O. Bucci and G. Franceschetti, "On the degrees of freedom of scattered fields," IEEE Trans. Antennas Propag. **37**, 918–926 (1989).
11. F. Gori, *Advanced Topics in Shannon Sampling and Interpolation Theory* (Springer-Verlag, 1993), pp. 37–83.
12. A. Lohmann, R. Dorsch, D. Mendlovic, Z. Zalevsky, and C. Ferreira, "Space-bandwidth product of optical signals and systems," J. Opt. Soc. Am. A **13**, 470–473 (1996).
13. D. Mendlovic and A. W. Lohmann, "Space-bandwidth product adaptation and its application to superresolution: fundamentals," J. Opt. Soc. Am. A **14**, 558–562 (1997).
14. R. Pierri and F. Soldovieri, "On the information content of the radiated fields in the near zone over bounded domains," Inverse Probl. **14**, 321–337 (1998).
15. R. Piestun and D. A. B. Miller, "Electromagnetic degrees of freedom of an optical system," J. Opt. Soc. Am. A **17**, 892–902 (2000).
16. D. A. B. Miller, "Communicating with waves between volumes: evaluating orthogonal spatial channels and limits on coupling strengths," Appl. Opt. **39**, 1681–1699 (2000).
17. J. Xu and R. Janaswamy, "Electromagnetic degrees of freedom in 2-d scattering environments," IEEE Trans. Antennas Propag. **54**, 3882–3894 (2006).
18. F. S. Oktem and H. M. Ozaktas, "Equivalence of linear canonical transform domains to fractional Fourier domains and the bicanonical width product: a generalization of the space–bandwidth product," J. Opt. Soc. Am. A **27**, 1885–1895 (2010).
19. F. S. Roux, "Complex-valued Fresnel-transform sampling," Appl. Opt. **34**, 3128–3135 (1995).
20. A. Stern and B. Javidi, "Analysis of practical sampling and reconstruction from Fresnel fields," Opt. Eng. **43**, 239–250 (2004).
21. L. Onural, "Exact analysis of the effects of sampling of the scalar diffraction field," J. Opt. Soc. Am. A **24**, 359–367 (2007).
22. J. J. Healy, B. M. Hennelly, and J. T. Sheridan, "Additional sampling criterion for the linear canonical transform," Opt. Lett. **33**, 2599–2601 (2008).
23. D. MacKay, "Quantal aspects of scientific information," Trans. IRE Prof. Group Inf. Theory **1**, 60–80 (1953).
24. J. T. Winthrop, "Propagation of structural information in optical wave fields," J. Opt. Soc. Am. **61**, 15–30 (1971).
25. T. W. Barret, "Structural information theory," J. Acoust. Soc. Am. **54**, 1092–1098 (1973).
26. H. M. Ozaktas, A. Koç, I. Sari, and M. A. Kutay, "Efficient computation of quadratic-phase integrals in optics," Opt. Lett. **31**, 35–37 (2006).
27. A. Koç, H. M. Ozaktas, C. Candan, and M. A. Kutay, "Digital computation of linear canonical transforms," IEEE Trans. Signal Process. **56**, 2383–2394 (2008).
28. F. Oktem and H. Ozaktas, "Exact relation between continuous and discrete linear canonical transforms," IEEE Signal Process. Lett. **16**, 727–730 (2009).
29. J. J. Healy and J. T. Sheridan, "Space–bandwidth ratio as a means of choosing between Fresnel and other linear canonical transform algorithms," J. Opt. Soc. Am. A **28**, 786–790 (2011).
30. F. T. Yu, *Optics and Information Theory* (Wiley, 1976).
31. M. J. Bastiaans, "Uncertainty principle and informational entropy for partially coherent light," J. Opt. Soc. Am. A **3**, 1243–1246 (1986).
32. M. S. Hughes, "Analysis of digitized wave-forms using Shannon entropy," J. Acoust. Soc. Am. **93**, 892–906 (1993).
33. D. Blacknell and C. J. Oliver, "Information-content of coherent images," J. Phys. D **26**, 1364–1370 (1993).
34. R. Barakat, "Some entropic aspects of optical diffraction imagery," Opt. Commun. **156**, 235–239 (1998).
35. M. A. Neifeld, "Information, resolution, and space-bandwidth product," Opt. Lett. **23**, 1477–1479 (1998).
36. F. T. Yu, *Entropy and Information Optics* (Marcel Dekker, 2000).
37. H. M. Ozaktas, S. Yüksel, and M. A. Kutay, "Linear algebraic theory of partial coherence: discrete fields and measures of partial coherence," J. Opt. Soc. Am. A **19**, 1563–1571 (2002).
38. A. Thaning, P. Martinsson, M. Karelin, and A. T. Friberg, "Limits of diffractive optics by communication modes," J. Opt. A **5**, 153–158 (2003).
39. A. Stern and B. Javidi, "Shannon number and information capacity of three-dimensional integral imaging," J. Opt. Soc. Am. A **21**, 1602–1612 (2004).
40. M. A. Porras and R. Medina, "Entropy-based definition of laser beam spot size," Appl. Opt. **34**, 8247–8251 (1995).
41. P. Réfrégier and J. Morio, "Shannon entropy of partially polarized and partially coherent light with Gaussian fluctuations," J. Opt. Soc. Am. A **23**, 3036–3044 (2006).
42. M. Migliore, "On electromagnetics and information theory," IEEE Trans. Antennas Propag. **56**, 3188–3200 (2008).

43. E. D. Micheli and G. A. Viano, "Inverse optical imaging viewed as a backward channel communication problem," J. Opt. Soc. Am. A **26**, 1393–1402 (2009).

44. A. Özçelikkale, H. M. Ozaktas, and E. Arıkan, "Signal recovery with cost constrained measurements," IEEE Trans. Signal Process. **58**, 3607–3617 (2010).

45. R. Konsbruck, E. Telatar, and M. Vetterli, "On sampling and coding for distributed acoustic sensing," IEEE Trans. Inf. Theory **58**, 3198–3214 (2012).

46. A. Özçelikkale and H. M. Ozaktas, "Representation of optical fields using finite numbers of bits," Opt. Lett. **37**, 2193–2195 (2012).

47. B. Dulek and S. Gezici, "Average Fisher information maximisation in presence of cost-constrained measurements," Electron. Lett. **47**, 654–656 (2011).

48. B. Dulek and S. Gezici, "Cost minimization of measurement devices under estimation accuracy constraints in the presence of Gaussian noise," Digit. Signal Process. **22**, 828–840 (2012).

49. H. L. Van Trees, *Detection, Estimation and Modulation Theory, Part I* (Wiley, 2001).

50. E. Collett and E. Wolf, "Beams generated by Gaussian quasi-homogeneous sources," Opt. Commun. **32**, 27–31 (1980).

51. F. Gori, "Collett-Wolf sources and multimode lasers," Opt. Commun. **34**, 301–305 (1980).

52. A. Starikov and E. Wolf, "Coherent-mode representation of Gaussian–Schell model sources and of their radiation fields," J. Opt. Soc. Am. A **72**, 923–928 (1982).

53. A. T. Friberg and R. J. Sudol, "Propagation parameters of Gaussian Schell-model beams," Opt. Commun. **41**, 383–387 (1982).

54. A. T. Friberg and J. Turunen, "Imaging of Gaussian–Schell model sources," J. Opt. Soc. Am. A **5**, 713–720 (1988).

55. P. Jixiong, "Waist location and Rayleigh range for Gaussian Schell-model beams," J. Opt. **22**, 157–159 (1991).

56. G. Gbur and E. Wolf, "The Rayleigh range of Gaussian Schell-model beams," J. Mod. Opt. **48**, 1735–1741 (2001).

57. L. Mandel and E. Wolf, *Optical Coherence and Quantum Optics* (Cambridge University, 1995).

58. Q. Lin and Y. Cai, "Fractional Fourier transform for partially coherent Gaussian–Schell model beams," Opt. Lett. **27**, 1672–1674 (2002).

59. S. Zhu, Y. Cai, and O. Korotkova, "Propagation factor of a stochastic electromagnetic Gaussian Schell-model beam," Opt. Express **18**, 12587–12598 (2010).

60. M. J. Bastiaans, "Wigner distribution function and its application to first-order optics," J. Opt. Soc. Am. A **69**, 1710–1716 (1979).

61. H. M. Ozaktas, Z. Zalevsky, and M. A. Kutay, *The Fractional Fourier Transform with Applications in Optics and Signal Processing* (Wiley, 2001).

62. K. B. Wolf, *Integral Transforms in Science and Engineering* (Plenum, 1979).

63. A. Balakrishnan, "A note on the sampling principle for continuous signals," IEEE Trans. Inf. Theory **3**, 143–146 (1957).

64. W. A. Gardner, "A sampling theorem for nonstationary random processes," IEEE Trans. Inf. Theory **18**, 808–809 (1972).

65. F. Garcia, I. Lourtie, and J. Buescu, "$L_2(R)$ nonstationary processes and the sampling theorem," IEEE Signal Process. Lett. **8**, 117–119 (2001).

66. W. H. Carter and E. Wolf, "Correlation theory of wavefields generated by fluctuating, three-dimensional, scalar sources," Opt. Acta **28**, 245–259 (1981).

67. R. A. Horn and C. R. Johnson, *Matrix Analysis* (Cambridge University, 1990).

68. T. M. Cover and J. A. Thomas, *Elements of Information Theory* (Wiley, 1991).

69. J. Buescu, "Positive integral operators in unbounded domains," J. Math. Anal. Appl. **296**, 244–255 (2004).