

Compressed Sensing Performance Bounds Under Poisson Noise

Maxim Raginsky, *Member, IEEE*, Rebecca M. Willett, *Member, IEEE*, Zachary T. Harmany, *Student Member, IEEE*, and Roummel F. Marcia, *Member, IEEE*

Abstract—This paper describes performance bounds for compressed sensing (CS) where the underlying sparse or compressible (sparsely approximable) signal is a vector of nonnegative intensities whose measurements are corrupted by Poisson noise. In this setting, standard CS techniques cannot be applied directly for several reasons. First, the usual signal-independent and/or bounded noise models do not apply to Poisson noise, which is nonadditive and signal-dependent. Second, the CS matrices typically considered are not feasible in real optical systems because they do not adhere to important constraints, such as nonnegativity and photon flux preservation. Third, the typical $\ell_2 - \ell_1$ minimization leads to overfitting in the high-intensity regions and oversmoothing in the low-intensity areas. In this paper, we describe how a feasible positivity- and flux-preserving sensing matrix can be constructed, and then analyze the performance of a CS reconstruction approach for Poisson data that minimizes an objective function consisting of a negative Poisson log likelihood term and a penalty term which measures signal sparsity. We show that, as the overall intensity of the underlying signal increases, an upper bound on the reconstruction error decays at an appropriate rate (depending on the compressibility of the signal), but that for a fixed signal intensity, the error bound actually grows with the number of measurements or sensors. This surprising fact is both proved theoretically and justified based on physical intuition.

Index Terms—Complexity regularization, compressive sampling, nonparametric estimation, photon-limited imaging, sparsity.

I. INTRODUCTION

THE basic idea of compressed sensing (CS) is that, when the signal of interest is very sparse (i.e., zero-valued at most locations) or highly compressible in some basis, relatively few “incoherent” observations are sufficient to reconstruct the most significant nonzero signal components [1], [2]. Despite the promise of this theory for many applications, very little is known

Manuscript received October 27, 2009; accepted April 06, 2010. Date of publication May 10, 2010; date of current version July 14, 2010. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Jerome Idler. This work was supported in part by an NSF CAREER Award No. CCF-06-43947, by a DARPA Grant No. HR0011-07-1-003, by an NGA Grant HM1582-10-1-0002, and by an NSF Grant DMS-08-11062. The material in this paper was presented at the IEEE International Symposium on Information Theory, Seoul, Korea, June/July 2009 and at the IEEE Workshop on Statistical Signal Processing, Cardiff, Wales, U.K., August/September 2009.

M. Raginsky, R. M. Willett, and Z. T. Harmany are with the Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708 USA (e-mail: m.raginsky@duke.edu; willett@duke.edu; zth@duke.edu).

R. F. Marcia is with the School of Natural Sciences, University of California, Merced, CA 95343 USA (e-mail: rmarcia@ucmerced.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2010.2049997

about its applicability to photon-limited imaging systems, where high-quality photomultiplier tubes are expensive and physically large, limiting the number of observations that can reasonably be acquired by the system. Limited photon counts arise in a wide variety of applications, including infrared imaging, nuclear medicine, astronomy and night vision, where the number of photons detected is very small relative to the number of pixels, voxels, or other entities to be estimated. Computational optics techniques, compressed sensing principles, and robust reconstruction methods can potentially lead to many novel imaging systems designed to make the best possible use of the small number of detected photons while reducing the size and cost of the detector array. Recent work has empirically explored CS in the context of photon limited measurements [3]–[6], but theoretical performance bounds similar to those widely cited in the conventional CS context previously remained elusive.

This is in part because the standard assumption of signal-independent and/or bounded noise (cf. [7] and [8]) is violated under the Poisson noise models used to describe images acquired by photon-counting devices [9]. The Poisson observation model

$$y \sim \text{Poisson}(Af^*) \quad (1)$$

where $f^* \in \mathbb{R}_+^m$ is the signal or image of interest, $A \in \mathbb{R}^{N \times m}$ linearly projects the scene onto an N -dimensional space of observations, and $y \in \mathbb{Z}_+^N$ is a length- N vector of observed Poisson counts, stipulates that the likelihood of observing a particular vector of counts y is given by

$$p(y|Af^*) = \prod_{j=1}^N \frac{(Af^*)_j^{y_j}}{y_j!} e^{-(Af^*)_j}$$

where $(Af^*)_j$ is the j th component of Af^* . Moreover, in order to correspond to a physically realizable linear optical system, the measurement matrix A must be:

- **Positivity-preserving**—for any nonnegative input signal f , the projected signal Af must also be nonnegative. Using the standard notation $f \succeq 0$ to denote the nonnegativity of f , we can write this condition as

$$f \succeq 0 \implies Af \succeq 0.$$

- **Flux-preserving**—for any input signal $f \succeq 0$, the mean total intensity of the observed signal Af must not exceed the total intensity incident upon the system:

$$\sum_{i=1}^N (Af)_i \leq \sum_{i=1}^m f_i.$$

A. Surprising Main Result

In this paper, we make the following contributions:

- design physically realizable sensing matrices, A , which incorporate the above positivity and photon flux preservation constraints;
- propose a penalized-likelihood objective function for reconstructing f^* from y observed according to (1);
- derive upper bounds on the error between f^* and the estimate \hat{f} and demonstrate how the error scales with the overall intensity ($I \triangleq \sum_i f_i^*$), the size of f^* (m), the number of measurements (N), and the compressibility of the signal in some basis (α); and
- present empirical results demonstrating the efficacy of the proposed method.

In particular, the main theoretical result presented in this paper shows that, for an α -compressible signal of total intensity I^1

$$\text{reconstruction error} \propto N \left(\frac{\log m}{I} \right)^{2\alpha/2\alpha+1} + \frac{\log \left(\frac{m}{N} \right)}{N}$$

for N sufficiently large. (As we show in Section III-B, there is a threshold effect in that the number of measurements N must be large enough to guarantee that the *per-sensor* reconstruction error decays with the incident signal intensity I .) Since the total number of observed events or photons, $n \triangleq \sum_{i=1}^N y_i$, is the realization of a Poisson random variable with intensity I , the bound reflects how error scales with the number of observed events.

While the rate of error decay as a function of the total intensity, I , coincides with earlier results in denoising contexts, the proportionality of the intensity-dependent term in the error to N may seem surprising at first glance. However, one can intuitively understand this result from the following perspective. If we increase the number of measurements (N) while keeping the expected number of observed photons (I) constant², the number of photons per sensor will decrease, so the signal-to-noise ratio (SNR) at each sensor will likewise decrease, thereby degrading performance. Having the number of sensors exceed the number of observed photons is not necessarily detrimental in a *denoising* or *direct measurement* setting (i.e., where A is the identity matrix) because multiscale algorithms can adaptively bin the noisy measurements together in homogeneous regions to achieve higher SNR overall [10], [11]. However, in the CS setting the signal is first altered by the compressive projections in the sensing matrix A , and the raw measurements cannot themselves be binned to improve SNR. In particular, there is no natural way to aggregate measurements across multiple sensors because the aggregation effectively changes the sensing matrix in a way that does not preserve critical properties of A .

¹More precisely, I refers to the total intensity integrated over the exposure time, so that increasing I can be associated with more source intensity, longer exposure time per measurement, or both.

²In some systems, such as a single-detector system, more measurements might seem to suggest more observed photons. However, holding I fixed while increasing N implies that each measurement is collected over a shorter exposure. Thus increasing N does *not* correspond to an increase in the number of observed events/photons.

One might also be surprised by this main result because in the case where the number of observed photons is very large (so that SNR is quite high and not a limiting factor), our bounds do not converge to the standard performance bounds in CS. This is because our bounds pertain to a sensing matrix A which, unlike conventional CS matrices based on i.i.d. realizations of a zero-mean random variable, is designed to correspond to a feasible physical system. In particular, every element of A must be nonnegative and appropriately scaled, so that the observed photon intensity is no greater than the photon intensity incident on the system (i.e., we cannot measure more light than is available). This rescaling dramatically impacts important elements of any performance bounds, including the form of the restricted isometry property [12], [13], even in the case of Gaussian or bounded noise. (Additional details and interpretation are provided in Section II-B after we introduce necessary concepts and notation.)

As a result, incorporating these real-world constraints into our measurement model has a *significant and adverse impact on the expected performance of an optical CS system*.

B. Relation to Previous CS Performance Bounds

The majority of the compressed sensing literature assumes that there exists a “sparsifying” reference basis W , so that $\theta^* \triangleq W^T f^*$ is sparse or lies in a weak- ℓ_p space. When the matrix product AW obeys the so-called *restricted isometry property* (RIP) [12], [13] or some related criterion, and when the noise is bounded or Gaussian, then θ^* can be accurately estimated from y by solving the following $\ell_2 - \ell_1$ optimization problem (or some variant thereof):

$$\hat{\theta} = \arg \min_{\theta} [\|y - AW\theta\|_2^2 + \tau\|\theta\|_1] \quad (2)$$

where $\tau > 0$ is a regularization parameter [2], [13], [14].

However, the ℓ_2 data-fitting term, $\|y - AW\theta\|_2^2$, is problematic in the presence of Poisson noise. Because under the Poisson model the variance of the noisy observations is proportional to the signal intensity, ℓ_2 data-fitting terms can lead to significant overfitting in high-intensity regions and oversmoothing in low-intensity regions. Furthermore, photon-limited imaging systems impose hard constraints on the nature of the measurements that can be collected, such as nonnegativity, which are not considered in much of the existing compressed sensing literature (recent papers of Dai and Milenkovic [15] and of Khajehnejad *et al.* [16] are notable exceptions). Bunea, Tsybakov, and Wegkamp [17] study the related problem of using ℓ_1 regularization for probability density estimation, but rather than assuming incoherent measurements of a random variable (similar to our setup), they assume direct observations of a random variable and learn, for example, a sparse mixture model. Furthermore, their work assumes infinite precision in the observed realizations of the random variable, so that their analysis does not provide any insight into how the number or bit depth of detector elements impacts performance. More recent work by Rish and Grabarnik [18] explores methods for CS reconstruction in the presence of exponential family noise using generalized linear models, but does not account for the physical constraints (such as flux preservation) associated with a typical Poisson setup.

In this paper, we propose estimating f^* from y using a regularized Poisson log-likelihood objective function as an alternative to (2), and we present risk bounds for recovery of a compressible signal from Poisson observations. Specifically, in the Poisson noise setting we maximize the log-likelihood while minimizing a penalty function that, for instance, could measure the sparsity of $\theta = W^T f$

$$\hat{f} = \arg \min_f \sum_{j=1}^N [(Af)_j - y_j \log(Af)_j] + \tau \text{pen}(f)$$

$$\text{subject to } f \succeq 0, \sum_{i=1}^m f_i = I \quad (3)$$

where $\text{pen}(\cdot)$ is a penalty function that will be detailed later, and I is the known total intensity of the unknown f^* .

C. Organization of the Paper

Section II contains the problem formulation, describes the proposed approach, and details the construction and properties of a feasible sensing matrix A . In Section III we develop an oracle inequality for the proposed estimator and then use it to establish risk bounds for compressible signals. Section IV contains a proof-of-concept experiment based on recent breakthroughs in sparse reconstruction methods we initially proposed in [4]. For the sake of readability, proofs of all theorems are relegated to the appendices.

II. PROBLEM FORMULATION AND PROPOSED APPROACH

We have a signal or image $f^* \succeq 0$ of size m that we wish to estimate using a detector array of size $N \ll m$. We assume that the total intensity of f^* , given by $I \triangleq \|f^*\|_1 = \sum_{i=1}^m f_i^*$, is known *a priori*. We make Poisson observations of Af^* , $y \sim \text{Poisson}(Af^*)$, where $A \in \mathbb{R}^{N \times m}$ is a positivity- and flux-preserving sensing matrix. Our goal is to estimate $f^* \in \mathbb{R}_+^m$ from $y \in \mathbb{Z}_+^N$.

A. Construction and Properties of the Sensing Matrix

We consider sensing matrices composed of zeros and (scaled) ones, where p is the probability of having a zero and $1 - p$ is the probability of having a one. In the context of optical systems, small p corresponds to sensing matrices with many ones, which allow most of the available light through the system to the detectors. Conversely, large p corresponds to having each measurement being the sum of a relatively small number of elements in the signal of interest. To explore the tradeoff between these two regimes, we explicitly consider p throughout our analysis.

We construct our sensing matrix A as follows. Given some $p \in (0, 1)$, let ν_p denote the probability distribution of a random variable that takes values

$$\lambda_p^- \triangleq -\sqrt{\frac{1-p}{p}} \text{ with probability } p;$$

$$\lambda_p^+ \triangleq \sqrt{\frac{p}{1-p}} \text{ with probability } 1-p.$$

Note that $\nu_{1/2}$ is the usual Rademacher distribution which puts equal mass on -1 and on $+1$. Let Z be an $N \times m$ matrix whose entries $Z_{i,j}$ are drawn i.i.d. from ν_p . We observe that

$$\mathbb{E}Z_{i,j} = 0 \text{ and } \mathbb{E}Z_{i,j}Z_{k,\ell} = \delta_{ik}\delta_{j\ell} \quad (4)$$

for all $1 \leq i, k \leq N$ and $1 \leq j, \ell \leq m$. Most compressed sensing approaches would proceed by assuming that we make (potentially noisy) observations of $\tilde{A}f^*$, where $\tilde{A} \triangleq Z/\sqrt{N}$. However, \tilde{A} will, with high probability, have at least one negative entry, which will render this observation model physically unrealizable in photon-counting systems. Therefore, we use A to generate a feasible sensing matrix A as follows. Let $\mathbf{1}_{r \times s}$ denote the $r \times s$ matrix all of whose entries are equal to 1. Then we take

$$A \triangleq \sqrt{\frac{p(1-p)}{N}} \tilde{A} + \frac{1-p}{N} \mathbf{1}_{N \times m}.$$

We can immediately deduce the following properties of A :

- It is positivity-preserving because each of its entries is either 0 or $1/N$.
- It is flux-preserving, i.e., for any $f \in \mathbb{R}_+^m$ we have

$$\|Af\|_1 \leq \|f\|_1. \quad (5)$$

This is easy to see directly: if $f \succeq 0$, then $Af \succeq 0$, and

$$\|Af\|_1 = \sum_{i=1}^N \sum_{j=1}^m A_{i,j} f_j \leq \sum_{j=1}^m f_j \equiv \|f\|_1.$$

- With probability at least $1 - Np^m$ (w.r.t. the realization of $\{Z_{i,j}\}$), every row of A has at least one nonzero entry. Assume that this event holds. Let $f \in \mathbb{R}^m$ be an arbitrary vector of intensities satisfying $f \succeq (cI)\mathbf{1}_{m \times 1}$ for some $c > 0$. Then

$$Af \succeq \frac{cI}{N} \mathbf{1}_{N \times 1}. \quad (6)$$

To see this, write

$$(Af)_i = \sum_{j=1}^m A_{i,j} f_j \geq \frac{1}{N} \cdot \min_{1 \leq j \leq m} f_j \geq \frac{cI}{N}.$$

Furthermore, and most importantly, with high probability \tilde{A} acts near-isometrically on certain subsets of \mathbb{R}^m . The usual formulation of this phenomenon is known in the compressed sensing literature as the restricted isometry property (RIP) [12], [13], where the subset of interest consists of all vectors with a given sparsity. However, as shown recently by Mendelson *et al.* [19], [20], the RIP is, in fact, a special case of a much broader circle of results concerning the behavior of random matrices whose entries are drawn from a *sub-Gaussian isotropic ensemble*. These terms are defined in Appendix VI, where we also prove the following result:

Theorem 1: Consider the matrix $\tilde{A} = Z/\sqrt{N}$, where the entries of Z are drawn i.i.d. from ν_p . Then there exist absolute constants $c_1, c_2 > 0$ such that the bound

$$\|u - v\|_2^2 \leq 4\|\tilde{A}u - \tilde{A}v\|_2^2 + \frac{2c_2^2\zeta_p^4 \log\left(\frac{c_2\zeta_p^4 m}{N}\right)}{N} \quad (7)$$

will hold simultaneously for all $u, v \in B_1^m$ with probability at least $1 - e^{-c_1 N/\zeta_p^4}$, where $B_1^m \triangleq \{u \in \mathbb{R}^m : \|u\|_1 = 1\}$ and

$$\zeta_p \triangleq \begin{cases} \sqrt{\frac{3}{2p(1-p)}}, & \text{if } p \neq \frac{1}{2} \\ 1, & \text{if } p = \frac{1}{2}. \end{cases} \quad (8)$$

Moreover, there exist absolute constants $c_3, c_4 > 0$ such that for any finite $T \subset S^{m-1}$, where $S^{m-1} \triangleq \{u \in \mathbb{R}^m : \|u\|_2 = 1\}$ is the unit sphere (ℓ_2)

$$\frac{1}{2} \leq \|\tilde{A}u\|_2^2 \leq \frac{3}{2}, \quad \forall u \in T \quad (9)$$

holds with probability at least $1 - e^{-c_3 N/\zeta_p^4}$, provided $N \geq c_4 \zeta_p^4 \log_2 |T|$.

B. dc Offset and Noise

The intensity underlying our Poisson observations can be expressed as

$$Af^* = \sqrt{\frac{p(1-p)}{N}} \tilde{A}f^* + \frac{I(1-p)}{N} \mathbf{1}_{N \times 1}.$$

As described in Theorem 1, the *idealized* sensing matrix \tilde{A} has a RIP-like property which can lead to certain performance guarantees if we could measure $\tilde{A}f^*$ directly; in this sense, $\tilde{A}f^*$ is the *informative* component of each measurement. However, a constant dc offset proportional to I is added to each element of $\tilde{A}f^*$ before Poisson measurements are collected, and elements of $\tilde{A}f^*$ will be very small relative to I . Thus the intensity and variance of each measurement will be proportional to I , overwhelming the informative elements of $\tilde{A}f^*$.³ (To see this, note that y_i can be approximated as $(Af^*)_i + \sqrt{(Af^*)_i} \xi_i$, where ξ_i is a Gaussian random variable with variance one.)

As we will show in this paper, the Poisson noise variance associated with the dc offset, necessary to model feasible measurement systems, leads to very different performance guarantees than are typically reported in the CS literature. The necessity of a dc offset is certainly not unique to our choice of a Rademacher sensing matrix; it has been used in practice for a wide variety of linear optical CS architectures (cf. [21]–[24]). A notable exception to the need for dc offsets is the expander-graph approach to generating nonnegative sensing matrices, which has recently been applied in Poisson CS settings [25]; however, theoretical results there are limited to signals which are sparse in the canonical (i.e., Dirac delta or pixel) basis.

As a result, the framework and bounds established in this paper have significant and sobering implications for any linear optical CS architecture operating in low-light settings.

³We would like to thank Emmanuel Candès and an anonymous reviewer for helpful insights on this point.

C. Reconstruction Approach and Bounds

We propose solving the following optimization problem:

$$\hat{f} \triangleq \arg \min_{f \in \Gamma} \left[-\log p(y|Af) + 2 \text{pen}(f) \right] \quad (10)$$

where $\text{pen}(f) \geq 0$ is a penalty term. Here, $\Gamma \equiv \Gamma(m, I)$ is a countable set of feasible estimators $f \in \mathbb{R}_+^m$ satisfying $\sum_{i=1}^m f_i = I$, and the penalty function satisfies the *Kraft inequality*

$$\sum_{f \in \Gamma} e^{-\text{pen}(f)} \leq 1. \quad (11)$$

[In (2) and (3), τ is a free parameter that could be selected by the user, while in (10) it is fixed at 2 for a penalty function that satisfies the Kraft inequality. In practice, one often prefers to use a value of τ different from what is supported in theory because of slack in the bounds.] While the penalty term may be chosen to be smaller for sparser solutions $\theta = W^T f$, where W is an orthogonal matrix that represents f in its “sparsifying” basis, our main result only assumes that (11) is satisfied. In fact, a variety of penalization techniques can be used in this framework; see [10], [26], and [27] for examples and discussions relating Kraft-compliant penalties to prefix codes for estimators. Many penalization or regularization methods in the literature, if scaled appropriately, can be considered prefix codelengths. We can think of (10) as a discretized-feasibility version of (3), where we optimize over a countable set of feasible vectors that grows in a controlled way with signal length m .

We will bound the accuracy with which we can estimate f^*/I ; in other words, we focus on accurately estimating the *distribution* of intensity in f^* independent of any scaling factor proportional to the total intensity of the scene, which is typically of primary importance to practitioners. Since the total number of observed events, n , obeys a Poisson distribution with mean I , estimating I by n is the strategy employed by most methods. However, the variance of this estimate is I , which means that, as I increases, our ability to estimate the distribution improves, while accurately estimating the unnormalized intensity is more challenging. We assume I is known to discount this effect. The quality of a candidate estimator f will be measured in terms of the *risk*

$$R(f^*, f) \triangleq \frac{1}{I^2} \|f^* - f\|_2^2.$$

D. Summary of Notation

Before proceeding to state and prove risk bounds for the proposed estimator, we summarize for the reader’s convenience the principal notation used in the sequel:

- m : dimension of the original signal
- $N (\ll m)$: number of measurements (detectors)
- $f^* \in \mathbb{R}_+^m$: unknown nonnegative-valued signal
- $I = \sum_i f_i^*$: total intensity of f^* , assumed known
- $Z \in \mathbb{R}^{N \times m}$: random matrix with i.i.d. entries taking values $-\sqrt{(1-p)/p}$ with probability p and $\sqrt{p/(1-p)}$ with probability $1-p$, where $p \in (0, 1)$ is a tunable parameter

- $\tilde{A} = Z/\sqrt{N}$: scaled matrix Z (cf. Theorem 1 for its norm preservation properties)
- ζ_p : sub-Gaussianity constant of \tilde{A} , defined in (8)
- c_1, c_2, c_3, c_4 : absolute constants from Theorem 1
- $A = \sqrt{p(1-p)/N}\tilde{A} + N^{-1}p(1-p)\mathbf{1}_{N \times m}$: physically realizable sensing matrix, obtained by shifting and scaling of \tilde{A} ; satisfies positivity and flux preservation requirements
- $\Gamma \subset \mathbb{R}_+^m$: finite or countably infinite set of candidate estimators with a penalty function $\text{pen} : \Gamma \rightarrow \mathbb{R}_+$ satisfying the Kraft inequality (11)
- $R(f^*, f) = \|f^* - f\|_2^2/I^2$: the risk of a candidate estimator f
- \hat{f} : the penalized maximum-likelihood estimator taking values in Γ , given by the solution to (10)

Other notation will be defined as needed in the appropriate sections.

III. BOUNDS ON THE EXPECTED RISK

Now we are in a position to establish risk bounds for the proposed estimator (10). Theorem 2 in Section III-A is a general risk bound that holds (with high probability w.r.t. the realization of \tilde{A}) for any sufficiently regular class of candidate estimators and a suitable penalty functional satisfying the Kraft inequality. Section III-B then particularizes Theorem 2 to the case in which the unknown signal f^* is compressible in some known reference basis, and the penalty is proportional to the sparsity of a candidate estimator in the reference basis.

A. An Oracle Inequality for the Expected Risk

In this section we give an upper bound on the expected risk $\mathbb{E}R(f^*, \hat{f})$ that holds for any target signal $f^* \succeq 0$ satisfying the normalization constraint $\sum_{i=1}^m f_i^* = I$, without assuming anything about the sparsity properties of f^* . Conceptually, our bound is an *oracle inequality*, which states that the expected risk of our estimator is within a constant factor of the best regularized risk attainable by estimators in Γ with full knowledge of the underlying signal f^* . More precisely, for each $f \in \Gamma$ define

$$R^*(f^*, f) \triangleq \frac{1}{I^2} \|f^* - f\|_2^2 + \frac{2 \text{pen}(f)}{I}$$

and for every $\Upsilon \subseteq \Gamma$ define

$$R^*(f^*, \Upsilon) \triangleq \min_{f \in \Upsilon} R^*(f^*, f)$$

i.e., the best penalized risk that can be attained over Υ by an oracle that has full knowledge of f^* . We then have the following.

Theorem 2: Suppose that, in addition to the conditions stated in Section II, the feasible set Γ also satisfies

$$Af \succeq \left(\frac{cI}{N}\right) \mathbf{1}_{N \times 1}, \quad \forall f \in \Gamma \quad (12)$$

for some $0 < c < 1$. Let $\mathcal{G}_{N,p}$ be the collection of all subsets $\Upsilon \subseteq \Gamma$, such that $|\Upsilon| \leq 2^{N/c_4 \zeta_p^4}$. Then the following holds

with probability at least $1 - me^{-KN}$ for some positive $K = K(c_1, c_3, p)$ (with respect to the realization of \tilde{A}):

$$\mathbb{E}R(f^*, \hat{f}) \leq C_{N,p} \min_{\Upsilon \in \mathcal{G}_{N,p}} R^*(f^*, \Upsilon) + \frac{2c_2^2 \zeta_p^4 \log\left(\frac{c_2 \zeta_p^4 m}{N}\right)}{N} \quad (13)$$

where

$$C_{N,p} \triangleq \max\left(\frac{24}{c}, \frac{16}{p(1-p)}\right) N$$

and the expectation is taken with respect to $y \sim \text{Poisson}(Af^*)$.

Remark 1: One way to satisfy the positivity condition (12) is to construct Γ so that

$$f \succeq (cI) \mathbf{1}_{m \times 1}, \quad \forall f \in \Gamma. \quad (14)$$

The desired condition (12) then follows from (6). A condition similar to (14) is natural in the context of estimating vectors with nonnegative entries from count data, as it excludes the possibility of assigning zero intensity to an input of a detector when at least one photon has been counted [28].

Remark 2: Both $C_{N,p}$ and ζ_p are minimized when $p = 1/2$, suggesting that altering the sensing architecture to have $p \neq 1/2$ may impair performance, despite the fact that increasing p would increase the expected total number of observed events (photons) and decreasing p would decrease the dc offset of the measurements and hence measurement noise variance.

Remark 3: Observe that for any pair $N_1 < N_2$ we have the inclusion $\mathcal{G}_{N_1,p} \subseteq \mathcal{G}_{N_2,p}$, which implies that $\min_{\Upsilon \in \mathcal{G}_{N,p}} R^*(f^*, \Upsilon)$ is a decreasing function of N . Hence, the first term on the right-hand side (RHS) of (13) is the product of a quantity that increases with N (i.e., $C_{N,p}$) and one that decreases with N . Combined with the presence of the $O(N^{-1} \log(m/N))$ additive term, this points to the possibility of a threshold effect, i.e., the existence of a critical number of measurements N^* , below which the expected risk may not monotonically decrease with N or I .

B. Risk Bounds for Compressible Signals

We now use Theorem 2 to analyze the performance of the proposed estimator when the target signal f^* is compressible (i.e., admits a sparse approximation) in some orthonormal reference basis.

Following [1], we assume that there exists an orthonormal basis $\Phi = \{\phi_1, \dots, \phi_m\}$ of \mathbb{R}^m , such that f^* is *compressible* in Φ in the following sense. Let W be the orthogonal matrix with columns ϕ_1, \dots, ϕ_m . Then the vector θ^* of the coefficients $\theta_j^* = \langle f^*, \phi_j \rangle$ of f^* in Φ is related to f^* via $f^* = W\theta^*$. Let $\theta_{(1)}^*, \dots, \theta_{(m)}^*$ be the decreasing rearrangement of θ^* : $|\theta_{(1)}^*| \geq |\theta_{(2)}^*| \geq \dots \geq |\theta_{(m)}^*|$. We assume that there exist some $0 < q < \infty$ and $\rho > 0$, such that

$$|\theta_{(j)}^*| \leq \rho I j^{-1/q}, \quad j = 1, \dots, m. \quad (15)$$

Note that for every $1 \leq j \leq m$ we have

$$|\theta_{(j)}^*| \leq \|\theta^*\|_2 = \|f^*\|_2 \leq \|f^*\|_1 = I$$

so we can take ρ to be a constant independent of I or m . Any θ^* satisfying (15) is said to belong to the *weak- ℓ_q ball of radius ρI* . The weak- ℓ_q condition (15) translates into the following approximation estimate [1]: given any $1 \leq k \leq m$, let $\theta^{(k)}$ denote the best k -term approximation to θ^* . Then

$$\frac{1}{I^2} \|\theta^* - \theta^{(k)}\|_2^2 \leq C\rho^2 k^{-2\alpha}, \quad \alpha = \frac{1}{q} - \frac{1}{2} \quad (16)$$

for some constant $C > 0$ that depends only on q . We also assume that f^* satisfies the condition (14) for some $c \in (0, 1)$, a lower bound on which is assumed known.

Theorem 3: There exist a finite set of candidate estimators Γ and a penalty function satisfying Kraft’s inequality, such that the bound

$$\mathbb{E}R(f^*, \hat{f}) \leq O(N) \min_{1 \leq k \leq k_*(N)} \left[k^{-2\alpha} + \frac{k}{m} + \frac{k \log_2 m}{I} \right] + O\left(\frac{\log\left(\frac{m}{N}\right)}{N}\right) \quad (17)$$

where

$$k_*(N) \triangleq \frac{N}{2c_4 \zeta_p^4 \log_2 m}$$

holds with the same probability as in Theorem 2. The constants obscured by the $O(\cdot)$ notation depend only on α, p, ρ, C and c .

The proof is presented in Appendix III; here we highlight a number of implications:

- 1) In the low-intensity setting $I \leq m \log m$, we get the risk bound

$$\mathbb{E}R(f^*, \hat{f}) \leq O(N) \min_{1 \leq k \leq k_*(N)} \left[k^{-2\alpha} + \frac{2k \log_2 m}{I} \right] + O\left(\frac{\log\left(\frac{m}{N}\right)}{N}\right).$$

If $k_*(N) \geq (\alpha I / \log_2 m)^{1/(2\alpha+1)}$, then we can further obtain

$$\mathbb{E}R(f^*, \hat{f}) \leq O(N) \left(\frac{I}{\log m}\right)^{-2\alpha/2\alpha+1} + O\left(\frac{\log\left(\frac{m}{N}\right)}{N}\right).$$

If $k_*(N) < (\alpha I / \log_2 m)^{1/(2\alpha+1)}$, there are not enough measurements, and the estimator saturates, although its risk can be controlled.

- 2) In the high-intensity setting $I > m \log m$, we obtain

$$\mathbb{E}R(f^*, \hat{f}) \leq O(N) \min_{1 \leq k \leq k_*(N)} \left[k^{-2\alpha} + \frac{2k}{m} \right] + O\left(\frac{\log\left(\frac{m}{N}\right)}{N}\right).$$

If $k_*(N) \geq (\alpha m)^{1/(2\alpha+1)}$, then we can further get

$$\mathbb{E}R(f^*, \hat{f}) \leq O(N) m^{-2\alpha/2\alpha+1} + O\left(\frac{\log\left(\frac{m}{N}\right)}{N}\right).$$

Again, if $k_*(N) < (\alpha m)^{1/(2\alpha+1)}$, there are not enough measurements, and the estimator saturates.

- 3) When $I \asymp m$ and $N \asymp m^{1/\beta}$ for some $\beta > 1 + 1/2\alpha$, we get (up to log terms) the rates

$$\mathbb{E}R(f^*, \hat{f}) = O(m^{-\gamma})$$

where $\gamma = [2\alpha - (2\alpha + 1)/\beta]/(2\alpha + 1) > 0$.

IV. EXPERIMENTAL RESULTS

In this section, we present the results of a proof-of-concept experiment showing the effectiveness of sparsity-regularized Poisson log likelihood reconstruction from CS measurements. In previous work [4], we described an optimization formulation called Sparse Poisson Intensity Reconstruction Algorithms (SPIRAL) for solving a simpler variant of (3)

$$\hat{f} = \arg \min_f [\phi(f) + \tau \text{pen}(f)] \text{ subject to } f \geq 0 \quad (18)$$

where $\phi(f) = \sum_j (Af)_j - y_j \log(Af)_j$. In our setup, since A has nonnegative entries, the constraint $Af \geq 0$ in (3) is redundant. Additionally, we do not require that the total intensity of the reconstruction f must sum to I since the resulting problem is easier to solve, and this equality constraint, in general, is approximately satisfied at the solution, i.e., it is not necessary to obtain accurate experimental results.

The proposed approach sequentially approximates the objective function in (18) by separable quadratic problems (QP) that are easier to minimize. In particular, at the k th iteration we use the second-order Taylor expansion of ϕ around f^k and approximate the Hessian by a positive scalar (η_k) multiple of the identity matrix, resulting in the following minimization problem:

$$f^{k+1} = \arg \min_f \left\| \left[f^k - \frac{1}{\eta_k} \nabla \phi(f^k) \right] - f \right\|_2^2 + \frac{2\tau}{\eta_k} \text{pen}(f) \text{ subject to } f \geq 0 \quad (19)$$

which can be viewed as a denoising subproblem applied to the gradient descent. This gives us considerable flexibility in designing effective penalty functions and in particular allows us to incorporate sophisticated “sparsity models” which encode, for instance, persistence of significant wavelet coefficients across scales to improve reconstruction performance. In the experiments below we utilize one such penalty, a partition-based estimator, as described in [10].

Partition-based methods calculate image estimates by determining the ideal partition of the domain of observations and by using maximum likelihood estimation to fit a model (e.g., a constant) to each interval in the optimal partition. The space of possible partitions is a nested hierarchy defined through a recursive dyadic partition (RDP) of the image domain, and the optimal partition is selected by pruning a quad-tree representation of the observed data to best fit the observations with minimal complexity. We call this partition-based algorithm SPIRAL-RDP. An additional averaging-over-shifts (cycle spinning) step can be efficiently incorporated to yield a translationally invariant (TI)

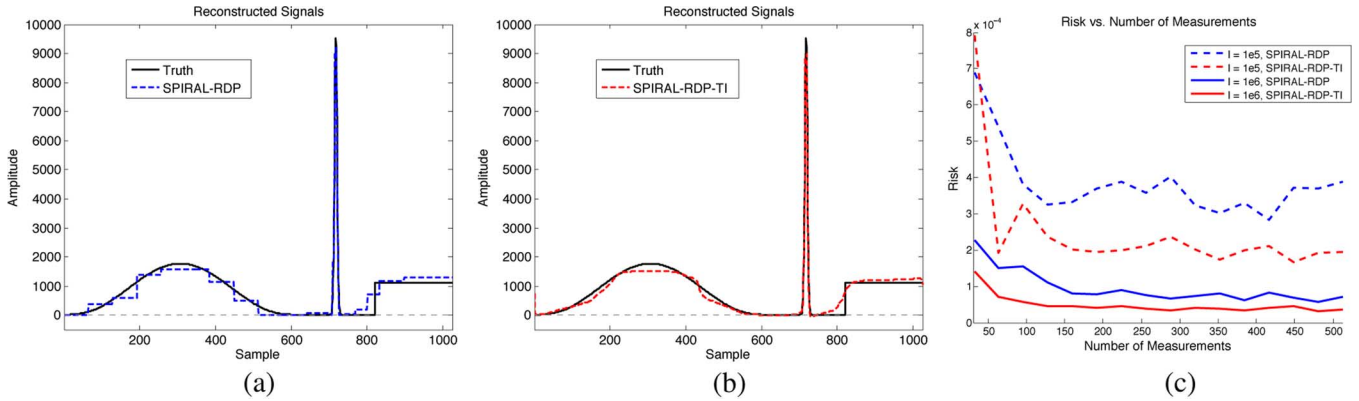


Fig. 1. Reconstruction results using the SPIRAL algorithm. (a) SPIRAL-RDP yields a nonsmooth piecewise constant estimate \hat{f}_{RDP} with a corresponding risk $R(\hat{f}_{\text{RDP}}, f^*) = 7.552 \times 10^{-5}$. (b) SPIRAL-RDP-TI produces a smoother and more accurate reconstruction $\hat{f}_{\text{RDP-TI}}$ with corresponding risk $R(\hat{f}_{\text{RDP-TI}}, f^*) = 4.468 \times 10^{-5}$. (c) shows the reconstruction accuracy versus number of measurements for two particular values of the intensity I (10^5 and 10^6).

algorithm, which we call SPIRAL-RDP-TI, that results in more accurate reconstructions.

The main computational costs of the SPIRAL methods come from matrix-vector multiplications involving A for calculating η_k and $\nabla \phi(x_k)$ in (19). Specifically, at each iteration k , SPIRAL computes two matrix-vector multiplications each with A and with A^T . For SPIRAL-RDP and SPIRAL-RDP-TI, even though the partition-based penalty QP appears difficult to solve because of its nonconvexity due to the penalty term, its global minimizer is easily computed using a noniterative tree-pruning algorithm (see [4] and [10] for details).

We evaluate the effectiveness of the proposed approaches in reconstructing a piecewise smooth function from noisy compressive measurements. In our simulations, the true signal [the black line in Fig. 1(a) and (b)] is of length 1024 and total intensity $I = 8.2 \times 10^5$. We take 512 noisy compressive measurements of the signal using a sensing matrix that contains 32 randomly distributed nonzero elements per row. This setup yields a mean detector photon count of 50, ranging from as few as 22 photons, to as high as 94 photons. We allowed each algorithm a fixed time budget of three seconds in which to run, which is sufficient to yield approximate convergence for all methods considered. Each algorithm was initialized at the same starting point: if $z = A^T y$, and $x : x_i = y_i / (Az)_i$, then we initialize with $f^0 : f_i^0 = z_i (A^T x)_i / (A^T \mathbf{1})_i$, where $\mathbf{1}$ is a vector of ones. The value of the regularization parameter τ was tuned independently for each algorithm to yield the minimal risk $R(\hat{f}, f^*) = \|\hat{f} - f^*\|_2^2 / I^2$ at the exhaustion of the computation budget. Tuning the regularization parameter in this manner is convenient in a simulation study. However, in the absence of truth, one typically resorts to a cross-validation procedure to determine an appropriate level of regularization.

In Fig. 1(a), we see that RDP estimates are less smooth than that of the original signal; however this drawback can be effectively neutralized through cycle spinning [see Fig. 1(b)]. In addition, the accuracy of the reconstruction (measured using the risk $R(\hat{f}, f^*)$) is improved by this averaging of shifts. Specifically, the SPIRAL-RDP estimate \hat{f}_{RDP} has a risk of $R(\hat{f}_{\text{RDP}}, f^*) = 7.552 \times 10^{-5}$, while the SPIRAL-RDP-TI estimate $\hat{f}_{\text{RDP-TI}}$

achieves a much lower risk of $R(\hat{f}_{\text{RDP-TI}}, f^*) = 4.468 \times 10^{-5}$. In Fig. 1(c), we examine how the performance of both partition-based SPIRAL methods scale with the number of measurements. These results utilize the same true signal and sensing matrix type as before, and are averaged over four noise realizations. By choosing two different intensity levels, we see that a higher intensity consistently leads to better performance. However, for a fixed intensity, the benefits of a higher number of measurements are less pronounced since collecting more observations necessarily results in a lower intensity per observation and hence noisier measurements (i.e., less photons collected per measurement). Note that the plot in Fig. 1(c) is not smoothly decreasing as one would expect; as we change the number of measurements, we need to randomly generate new Poisson realizations of our data, and hence there is some degree of variability in these results.

In [4], we examine our SPIRAL approach with an ℓ_1 -norm penalty on the coefficients of a wavelet expansion of the signal. In this case, the resulting reconstruction is very oscillatory with pronounced wavelet artifacts. With an increase in the regularization parameter these artifacts can be minimized; however, this leads to an “oversmoothed” solution and increases the risk of the estimate. In addition, we compare the SPIRAL approaches to the more established expectation-maximization algorithms based upon the same maximum penalized likelihood estimation objective function found in (19) and demonstrate that reconstructions from the partition-based SPIRAL methods are more accurate. In particular, they produce estimates with lower risk values, are more effective in recovering regions of low intensity, and yield reconstructions without spurious wavelet artifacts.

We mention other recent approaches for solving Poisson inverse problems; given that a detailed comparison of the performances of these various methods is beyond the scope of this paper, we simply outline some critical features with these approaches. In [29], the Poisson statistical model is bypassed in favor of an additive Gaussian noise model through the use of the Anscombe variance-stabilizing transform. This statistical simplification is not without cost, as the linear projections of the scene must now be characterized as nonlinear observations.

Other recent efforts [5], [30] solve Poisson image reconstruction problems with total variation norm regularization, but the method involves a matrix inverse operation, which can be extremely difficult to compute for large problems outside of deconvolution settings. Finally, the approaches in [31] and [32] apply proximal functions to solve more general constrained convex minimization problems. These methods use projection to obtain feasible iterates (i.e., nonnegative intensity values), which may be difficult for recovering signals that are sparse in a noncanonical basis.

V. CONCLUSION

We have derived upper bounds on the compressed sensing estimation error under Poisson noise for sparse or compressible signals. We specifically prove error decay rates for the case where the penalty term is proportional to the ℓ_0 -norm of the solution; this form of penalty has been used effectively in practice with a computationally efficient Expectation-Maximization algorithm (cf. [22]), but was lacking the theoretical support provided by this paper. Furthermore, the main theoretical result of this paper holds for any penalization scheme satisfying the Kraft inequality, and hence can be used to assess the performance of a variety of potential reconstruction strategies besides sparsity-promoting reconstructions.

One significant aspect of the bounds derived in this paper is that their signal-dependent portion grows with N , the size of the measurement array, which is a major departure from similar bounds in the Gaussian or bounded-noise settings. It does not appear that this is a simple artifact of our analysis. Rather, this behavior can be intuitively understood to reflect that elements of y will all have similar values at low light levels, making it very difficult to infer the relatively small variations in $\hat{A}f^*$. Hence, compressed sensing using shifted Rademacher sensing matrices is fundamentally difficult when the data are Poisson observations.

APPENDIX I

PROOF OF THEOREM 1: The proof makes heavy use of the geometric approach of [19] and [20]. Since this approach is not as well known in the compressed sensing community as the usual RIP, we give a brief exposition of its main tenets. Consider the problem of recovering an unknown signal f^* , which resides in some set $\Lambda \subset \mathbb{R}^m$, from N linear measurements of the form $y_1 = \langle Z_1, f^* \rangle, \dots, y_N = \langle Z_N, f^* \rangle$, where the measurement vectors $Z_1, \dots, Z_N \in \mathbb{R}^m$ are i.i.d. samples from a distribution μ which is:

- *sub-Gaussian with constant ζ* , i.e., there exists a constant $\zeta > 0$, such that for $Z_0 \sim \mu$ and for every $u \in \mathbb{R}^m$

$$\inf \left\{ s \geq 0 : \mathbb{E} \exp \left(\frac{|\langle Z_0, u \rangle|^2}{s^2} \right) \leq 2 \right\} \leq \zeta \|u\|_2 \quad (20)$$

- *isotropic*, i.e., for $Z_0 \sim \mu$ and for every $u \in \mathbb{R}^m$, $\mathbb{E} |\langle Z_0, u \rangle|^2 = \|u\|_2^2$.

The main results of [19] revolve around the norm preservation properties of the *random operator* $\hat{A} : \mathbb{R}^m \rightarrow \mathbb{R}^N$ defined by

$$\tilde{A}u \triangleq \frac{1}{\sqrt{N}} \sum_{i=1}^N \langle Z_i, u \rangle e_i \quad (21)$$

where e_1, \dots, e_N is the standard basis in \mathbb{R}^N . Particularized to the case of $\Lambda = B_1^m$, the unit ball w.r.t. the ℓ_1 norm on \mathbb{R}^m , the first main result of [19] reads as follows.

Theorem 4: [19, Theorem A] Let $Z_1, \dots, Z_N \in \mathbb{R}^m$ be i.i.d. samples from an ensemble μ which is isotropic and sub-Gaussian with constant $\zeta \geq 1$. There exist absolute constants $c_1, c_2 > 0$, such that, with probability at least $1 - e^{-c_1 N/\zeta^4}$

$$\|u - v\|_2^2 \leq 4 \|\tilde{A}(u - v)\|_2^2 + \frac{2c_2^2 \zeta^4 \log \left(\frac{c_2 \zeta^4 m}{N} \right)}{N} \quad (22)$$

for all $u, v \in B_1^m$.

The other main result of [19], informally, states the following: for any finite set $T \subset S^{m-1}$, the random operator \tilde{A} does not distort the norms of the elements of T too much, provided the number of measurements N is sufficiently large. The required minimal number of measurements is determined by the cardinality of T . In its precise form, the relevant result of [19] says the following.

Theorem 5: [19, Corollary 2.7] There exist absolute constants $c_3, c_4 > 0$, such that the following holds. Consider a finite set $T \subset S^{m-1}$, and let $Z_1, \dots, Z_N \in \mathbb{R}^m$ be i.i.d. samples from a ζ -sub-Gaussian isotropic ensemble. Then, with probability at least $1 - e^{-c_3 N/\zeta^4}$

$$\frac{1}{2} \leq \|\tilde{A}u\|_2^2 \equiv \frac{1}{N} \sum_{i=1}^N |\langle Z_i, u \rangle|^2 \leq \frac{3}{2}, \quad \forall u \in \Lambda \quad (23)$$

provided $N \geq c_4 \zeta^4 \log_2 |\Lambda|$.

Remark 4: Theorem 5 is a special case of a more general result that applies to *general* (not necessarily finite) subsets T of the unit sphere. The minimum necessary number of measurements is determined by the so-called ℓ_* -functional of T , which is defined as follows. Let g_1, \dots, g_m be independent standard Gaussian random variables, i.e., each $g_i \sim N(0, 1)$ independently of all others. Then

$$\ell_*(T) \triangleq \mathbb{E} \sup_{u \in T} \left| \sum_{i=1}^m g_i u_i \right|,$$

where the expectation is taken w.r.t. g_1, \dots, g_m , and u_i denotes the i th component of u . Informally, $\ell_*(T)$ measures how much the elements of T are correlated with white Gaussian noise. Estimates of $\ell_*(T)$ are available for many sets T . For instance (see [19, Section III] and references therein):

- If T is a finite set, then $\ell_*(T) \leq c \sqrt{\log_2 |T|}$ for some absolute constant $c > 0$.

- If T is the set of all convex combinations of unit vectors in \mathbb{R}^m whose ℓ_0 norm is at most k ,

$$T = \text{conv hull} \{u \in S^{m-1} : \|u\|_0 = |\{i : u_i \neq 0\}| \leq k\} \quad (24)$$

then $\ell_*(T) \leq c\sqrt{k \log_2(cm/k)}$ for some absolute constant $c > 0$. (We do not use this particular T in our analysis, but mention it here because of its connection to the more widely known RIP [12].)

The minimum necessary number of measurements for (9) to hold with the same probability as before is determined by $N \geq c_4 \zeta^4 \ell_*(T)^2$. When $|T|$ is finite, combining this bound with the estimate of $\ell_*(T)$ in terms of the log-cardinality of T yields Theorem 5. Moreover, as shown in [20], the usual RIP for matrices with rows drawn from sub-Gaussian isotropic ensembles is a consequence of this result as well. Specifically, it relies on the $\ell_*(T)$ estimate for the set T defined in (24).

We now apply Theorems 4 and 5 to the measurement matrix \tilde{A} defined in Section II-A. Recall that $\tilde{A} = Z/\sqrt{N}$, and let $Z_i = (Z_{i,1}, \dots, Z_{i,m})$ denote the i th row of Z . By construction, each Z_i is an independent copy of a random variable $Z_0 \in \mathbb{R}^m$ with distribution $\nu_p^{\otimes m}$ —that is, the components of Z_0 are drawn i.i.d. from ν_p . To be able to apply Theorems 4 and 5, we first show that the distribution $\nu_p^{\otimes m}$ is sub-Gaussian and isotropic. To that end, we need to obtain a bound of the form (20) for linear functionals of the form $\langle Z_0, u \rangle$. The infimum on the left-hand side (LHS) of (20) is the so-called *Orlicz ψ_2 -norm* of the random variable $\langle Z_0, u \rangle$. Given a real-valued random variable U , its Orlicz ψ_2 -norm [33, Ch. 2] is defined as

$$\|U\|_{\psi_2} \triangleq \inf \left\{ s \geq 0 : \mathbb{E} \exp \left(\frac{U^2}{s^2} \right) \leq 2 \right\}.$$

Thus, μ is sub-Gaussian with constant ζ if for $Z_0 \sim \mu$ we have

$$\|\langle Z_0, u \rangle\|_{\psi_2} \leq \zeta \|u\|_2, \quad \forall u \in \mathbb{R}^m.$$

Here is a useful tool for bounding Orlicz norms.

Lemma 1: [33, Lemma 2.2.1] Let U be a real-valued random variable that satisfies the tail bound

$$\mathbb{P}[|U| > t] \leq K e^{-Ct^2}$$

for all $t > 0$, where $K, C > 0$ are some constants. Then its Orlicz ψ_2 -norm satisfies $\|U\|_{\psi_2} \leq \sqrt{(1+K)/C}$.

Using this lemma, we can prove the following.

Lemma 2: The product probability measure $\nu_p^{\otimes m}$ is isotropic and sub-Gaussian with constant ζ_p defined in (8).

Proof: Let $Z_0 = (Z_{0,1}, \dots, Z_{0,m}) \sim \nu_p^{\otimes m}$. Isotropy is immediate from (4). To prove sub-Gaussianity, let us first assume that $p \neq 1/2$. Fix some $u \in \mathbb{R}^m$ and consider the random variable $\langle Z_0, u \rangle$, which is a sum of independent random variables $Z_{0,j}u_j$, $1 \leq j \leq m$. Then $\mathbb{E}\langle Z_0, u \rangle = 0$, and with probability one each $Z_{0,j}u_j$ takes values in the set $\{-\lambda_p^- |u_j|, -\lambda_p^+ |u_j|\}$ if $u_j < 0$, and in $\{\lambda_p^- |u_j|, \lambda_p^+ |u_j|\}$ if $u_j \geq 0$. Hence, Hoeffding's inequality gives the tail bound

$$\mathbb{P} \left[|\langle Z_0, u \rangle| > t \right] \leq 2 \exp \left(- \frac{2t^2}{(\lambda_p^+ - \lambda_p^-)^2 \sum_{j=1}^m |u_j|^2} \right)$$

$$= 2 \exp \left(- \frac{2p(1-p)t^2}{\|u\|_2^2} \right).$$

Using Lemma 1 with $K = 2$ and $C = 2p(1-p)/\|u\|_2^2$, we get the desired result. When $p = 1/2$, using the fact that the Rademacher measure is symmetric, it can be shown that $\nu_{1/2}^{\otimes m}$ is sub-Gaussian with constant $\zeta = 1$ [19]. ■

Now, using Theorems 4 and 5 in conjunction with Lemma 2, we have proved Theorem 1.

APPENDIX II

PROOF OF THEOREM 2: With probability at least $1 - e^{-c_3 N/\zeta_p^4}$, the following chain of estimates holds:

$$\begin{aligned} \frac{1}{I^2} \|f^* - \hat{f}\|_2^2 &\leq \frac{4}{I^2} \|\tilde{A}(f^* - \hat{f})\|_2^2 \\ &\quad + \frac{2c_2^2 \zeta_p^4 \log(c_2 \zeta_p^4 \frac{m}{N})}{N} \\ &= \frac{4N}{p(1-p)I^2} \|A(f^* - \hat{f})\|_2^2 \\ &\quad + \frac{2c_2^2 \zeta_p^4 \log(c_2 \zeta_p^4 \frac{m}{N})}{N} \\ &\leq \frac{4N}{p(1-p)I^2} \|A(f^* - \hat{f})\|_1^2 \\ &\quad + \frac{2c_2^2 \zeta_p^4 \log(c_2 \zeta_p^4 \frac{m}{N})}{N} \end{aligned}$$

where the first inequality is a consequence of the first part of Theorem 1, and the remaining steps follow from definitions, standard inequalities for ℓ_p norms, and the fact that $\sum_i f_i = \sum_i f_i^* = I$ for all $f \in \Gamma^4$. Moreover

$$\begin{aligned} \|A(f^* - \hat{f})\|_1^2 &= \left(\sum_{i=1}^N \left| (Af^*)_i^{1/2} - (A\hat{f})_i^{1/2} \right| \right. \\ &\quad \left. \cdot \left| (Af^*)_i^{1/2} + (A\hat{f})_i^{1/2} \right| \right)^2 \\ &\leq \sum_{i,j=1}^N \left| (Af^*)_i^{1/2} - (A\hat{f})_i^{1/2} \right|^2 \\ &\quad \cdot \left| (Af^*)_j^{1/2} + (A\hat{f})_j^{1/2} \right|^2 \\ &\leq 2 \sum_{i,j=1}^N \left| (Af^*)_i^{1/2} - (A\hat{f})_i^{1/2} \right|^2 \\ &\quad \cdot \left((Af^*)_j + (A\hat{f})_j \right) \\ &\leq 4I \sum_{i=1}^N \left| (Af^*)_i^{1/2} - (A\hat{f})_i^{1/2} \right|^2 \end{aligned}$$

where the first inequality follows from Cauchy-Schwarz, the second is a consequence of the arithmetic-mean/geometric-mean inequality, and the third follows from (5). It is a

⁴We use the fact that $\|A(f^* - \hat{f})\|_2^2 = N/p(1-p)\|\tilde{A}(f^* - \hat{f})\|_2^2$, which is only true when $I = \|f^*\|_1 = \|\hat{f}\|_1$. If we did not assume I was known, we would have additional terms in our error bound which would be proportional to I and would not reflect our ability to exploit prior knowledge of structure or sparsity to achieve an accurate solution.

matter of straightforward algebra (see Appendix IV) to show that

$$\begin{aligned} \sum_{i=1}^N \left| (Af^*)_i^{1/2} - (A\hat{f})_i^{1/2} \right|^2 &= -2 \log \prod_{i=1}^N \exp \\ &\quad \left(-\frac{1}{2} \left[(Af^*)_i^{1/2} - (A\hat{f})_i^{1/2} \right] \right)^2 \\ &= 2 \log \frac{1}{\int \sqrt{p(y|Af^*)p(y|A\hat{f})} d\nu(y)} \end{aligned} \quad (25)$$

where ν is the counting measure on \mathbb{Z}_+^N . Now, the same techniques as in Li and Barron [34] (see also the proof in [35, Th. 7] or Appendix IV-B) can be used to show that

$$2\mathbb{E} \log \frac{1}{\int \sqrt{p(y|Af^*)p(y|A\hat{f})} d\nu(y)} \leq \min_{f \in \Gamma} \left[\text{KL}(p(\cdot|Af^*)||p(\cdot|Af)) + 2 \text{pen}(f) \right] \quad (26)$$

where $\text{KL}(\cdot||\cdot)$ is the *Kullback-Leibler (KL) divergence*, which for the Poisson likelihood has the form

$$\begin{aligned} \text{KL}(p(\cdot|Af^*)||p(\cdot|Af)) &= \sum_{i=1}^N \left[(Af^*)_i \log \frac{(Af^*)_i}{(Af)_i} - (Af^*)_i + (Af)_i \right]. \end{aligned}$$

Using $\log t \leq t - 1$ together with (12), which holds with probability at least $1 - Np^m$, and (6), we can bound the KL divergence as

$$\begin{aligned} &\sum_{i=1}^N \left[(Af^*)_i \log \frac{(Af^*)_i}{(Af)_i} - (Af^*)_i + (Af)_i \right] \\ &\leq \sum_{i=1}^N \left[(Af^*)_i \left(\frac{(Af^*)_i}{(Af)_i} - 1 \right) - (Af^*)_i + (Af)_i \right] \\ &= \sum_{i=1}^N \frac{1}{(Af)_i} \left[(Af)_i^2 - 2(Af)_i(Af^*)_i + (Af^*)_i^2 \right] \\ &\leq \frac{N}{cI} \|A(f^* - f)\|_2^2 \\ &= \frac{p(1-p)}{cI} \|\tilde{A}(f^* - f)\|_2^2. \end{aligned}$$

Now, choose any $\Upsilon^* \in \mathcal{G}_{N,p}$, such that

$$R^*(f^*, \Upsilon^*) = \min_{\Upsilon \in \mathcal{G}_{N,p}} R^*(f^*, \Upsilon).$$

Observe that $N \geq c_4 \zeta_p^4 \log_2 |\Upsilon^*|$, so we can apply the second part of Theorem 1 to the set $\{(f^* - f)/\|f^* - f\|_2 : f \in \Upsilon^*\}$. Then, with probability at least $1 - e^{-c_3 N/\zeta_p^4}$, we have

$$\|\tilde{A}(f^* - f)\|_2^2 \leq \left(\frac{3}{2}\right) \|f^* - f\|_2^2, \quad \forall f \in \Upsilon^*.$$

Combining everything and rearranging, we get the bound

$$\mathbb{E}R(f^*, \hat{f}) \leq C_{N,p} \min_{f \in \Upsilon^*} \left[\left\| \frac{f^*}{I} - \frac{f}{I} \right\|_2^2 + \frac{2 \text{pen}(f)}{I} \right] + \frac{2c_2^2 \log\left(\frac{c_2 \zeta_p^4 m}{N}\right)}{N}$$

which holds with probability at least $1 - (Np^m + e^{-c_1 N/\zeta_p^4} + e^{-c_3 N/\zeta_p^4}) \geq 1 - me^{-KN}$ for a suitable constant $K = K(c_1, c_3, p)$. The theorem is proved.

APPENDIX III

PROOF OF THEOREM 3: In order to apply Theorem 2, we will form a suitable finite class of estimators Γ and set a penalty function $\text{pen}(f)$ over this class which (a) is smaller for sparser $\theta = W^T f$ and (b) satisfies (11). The family Γ is constructed as follows.

1) Define the sets

$$\begin{aligned} \Theta &\triangleq \{\theta \in \mathbb{R}^m : \|\theta\|_\infty \leq I; \text{ each } \theta_i \\ &\quad \text{uniformly quantized into one of} \\ &\quad \sqrt{m} \text{ bins over } [-I, I]\} \end{aligned}$$

and $\mathcal{F} \triangleq \{f = W\theta : \theta \in \Theta\}$.

2) For each $f \in \mathcal{F}$, let \bar{f} denote its ℓ_2 projection onto the closed convex set

$$\mathcal{C} \triangleq \left\{ g \in \mathbb{R}^m : g \succeq (cI)\mathbf{1}_{m \times 1} \text{ and } \sum_{i=1}^m g_i = I \right\}$$

i.e.

$$\bar{f} \triangleq \arg \min_{g \in \mathcal{C}} \|f - g\|_2.$$

3) Finally, let $\Gamma \triangleq \{\bar{\theta} = W^T \bar{f} : f \in \mathcal{F}\}$.

Note that the projection \bar{f} satisfies the *Pythagorean identity*

$$\|g - f\|_2^2 \geq \|g - \bar{f}\|_2^2 + \|\bar{f} - f\|_2^2, \quad \forall g \in \mathcal{C}$$

(see, e.g., [36, Th. 2.4.1]). In particular, $\|g - f\|_2^2 \geq \|g - \bar{f}\|_2^2$, and, since $f^* \in \mathcal{C}$, we have

$$\|f^* - \bar{f}\|_2^2 \leq \|f^* - f\|_2^2, \quad \forall f \in \mathcal{F}. \quad (27)$$

Consider the penalty

$$\text{pen}(f) = \log_2(m+1) + \left(\frac{3}{2}\right) \|\theta\|_0 \log_2(m), \quad \theta = W^T f$$

(measured in bits; in order to satisfy Kraft's inequality as stated in (11), it will need to be rescaled by $\log 2$). This corresponds to the following prefix code for $\theta \in \Theta$ (that is, we encode the elements of Θ , before they are subjected to the deterministic operation of projecting onto \mathcal{C}):

- 1) First we encode $\|\theta\|_0$, the number of nonzero components of θ , which can be encoded with $\log_2(m+1)$ bits.
- 2) For each of the $\|\theta\|_0$ nonzero components, we encode its location in the θ vector; since there are m possible locations, this takes $\log_2(m)$ bits per component.
- 3) Next we encode each coefficient value, quantized to one of \sqrt{m} uniformly sized bins.

Since this corresponds to a uniquely decodable code for $f \in \mathcal{F}$ (or $\theta \in \Theta$), we see that $\text{pen}(f)$ indeed satisfies the Kraft inequality [37].

Now, given $\theta^* = W^T f^*$, let $\theta^{(k)}$ be its best k -term approximation, $\theta_q^{(k)} \in \Theta$ the quantized version of $\theta^{(k)}$, for which we have

$$\frac{1}{I^2} \|\theta_q^{(k)} - \theta^{(k)}\|_2^2 \leq \frac{k}{m}$$

and $\bar{\theta}_q^{(k)}$ the element of Γ obtained by projecting $f_q^{(k)} = W\theta_q^{(k)}$ onto \mathcal{C} and then transforming back into the basis Φ : $\bar{\theta}_q^{(k)} = W^T \bar{f}_q^{(k)}$. Then, using (27) and (16), we get

$$\begin{aligned} \|f^* - \bar{f}_q^{(k)}\|_2^2 &\leq \|f^* - f_q^{(k)}\|_2^2 \\ &= \|\theta^* - \theta_q^{(k)}\|_2^2 \\ &\leq 2\|\theta^* - \theta^{(k)}\|_2^2 + 2\|\theta^{(k)} - \theta_q^{(k)}\|_2^2 \\ &\leq I^2 \left(2C\rho^2 k^{-2\alpha} + \frac{2k}{m} \right). \end{aligned}$$

Given each $1 \leq k \leq m$, let $\Gamma_k \subseteq \Gamma$ be the set of all $\bar{\theta} \in \Gamma$, such that the corresponding $\theta \in \Theta$ satisfies $|\theta|_0 \leq k$. Then $|\Gamma_k| = \binom{m}{k} m^{k/2}$, so that $\log_2 |\Gamma_k| \leq 2k \log_2 m$, and therefore $\Gamma_k \in \mathcal{G}_{N,p}$ whenever

$$k \leq k_*(N), \text{ where } k_*(N) \triangleq \frac{N}{2c_4 \zeta_p^4 \log_2 m}.$$

Then the first term on the RHS of (13) can be bounded by

$$\begin{aligned} C_N \min_{1 \leq k \leq k_*(N)} R^*(f^*, \Gamma_k) &\leq O(N) \min_{1 \leq k \leq k_*(N)} \left[\frac{1}{I^2} \|\theta^* - \bar{\theta}_q^{(k)}\|_2^2 \right. \\ &\quad \left. + \frac{2 \text{pen}(f_q^{(k)})}{I} \right] \\ &\leq O(N) \min_{1 \leq k \leq k_*(N)} \left[k^{-2\alpha} + \frac{k}{m} + \frac{k \log_2 m}{I} \right] \end{aligned}$$

where the constant obscured by the $O(\cdot)$ notation depends only on p, ρ, C and c .

APPENDIX IV

AUXILIARY TECHNICAL RESULTS:

Proof of (25): Given two Poisson intensity vectors $g, h \in \mathbb{R}_+^N$, we have

$$\begin{aligned} \int \sqrt{p(y|g)p(y|h)} d\nu(y) &= \prod_{i=1}^N \int \sqrt{p(y_i|g_i)p(y_i|h_i)} d\nu_i(y_i) \\ &= \prod_{i=1}^N \sum_{y_i=0}^{\infty} \frac{(g_i h_i)^{y_i/2}}{y_i!} e^{-(g_i+h_i)/2} \\ &= \prod_{i=1}^N e^{-(g_i-2(g_i h_i)^{1/2}+h_i)/2} \\ &\quad \sum_{y_i=0}^{\infty} \frac{(g_i h_i)^{y_i/2}}{y_i!} e^{-(g_i h_i)^{1/2}} \\ &= \prod_{i=1}^N e^{-1/2((g_i)^{1/2}-(h_i)^{1/2})^2} \\ &\quad \underbrace{\int p(y_i|(g_i h_i)^{1/2}) d\nu_i(y_i)}_{=1} \\ &= \prod_{i=1}^N e^{-1/2((g_i)^{1/2}-(h_i)^{1/2})^2} \end{aligned}$$

where ν_i denotes the counting measure on the i th component of y . Taking logs, we obtain

$$2 \log \frac{1}{\int \sqrt{p(y|g)p(y|h)} d\nu(y)} = \sum_{i=1}^N \left((g_i)^{1/2} - (h_i)^{1/2} \right)^2.$$

The quantity on the LHS is often used to measure divergence between probability distributions, and dates back to the work of Bhattacharyya [38] and Chernoff [39].

Proof of (26): For the sake of brevity, we will write $p_{f^*}(y)$ and $p_f(y)$ instead of $p(y|A_{f^*})$ and $p(y|A_f)$. Also, define the *Hellinger affinity*

$$\mathcal{A}(f^*, f) \triangleq \int \sqrt{p_{f^*}(y)p_f(y)} d\nu(y).$$

We then have

$$\begin{aligned} 2 \log \frac{1}{\mathcal{A}(f^*, \hat{f})} &= 2 \log \left[\frac{\sqrt{\frac{p_{\hat{f}}(y)}{p_{f^*}(y)} e^{-\text{pen}(\hat{f})}}{\mathcal{A}(f^*, \hat{f})}} \right] \\ &\quad + \log \frac{p_{f^*}(y)}{p_{\hat{f}}(y)} + 2 \text{pen}(\hat{f}). \end{aligned}$$

In the first term on the RHS, the ratio is evaluated at \hat{f} . Replacing this ratio by the sum of such ratios evaluated at every $f \in \Gamma$, we obtain the upper bound

$$2 \log \sum_{f \in \Gamma} \left[\frac{\sqrt{\frac{p_f(y)}{p_{f^*}(y)} e^{-\text{pen}(f)}}{\mathcal{A}(f^*, f)}} \right] + \log \frac{p_{f^*}(y)}{p_{\hat{f}}(y)} + 2 \text{pen}(\hat{f}).$$

Now we take expectation w.r.t. $p_{f^*}(y)$. Then, by Jensen's inequality

$$\begin{aligned} & \mathbb{E}_{f^*} \left\{ \log \sum_{f \in \Gamma} \left[\frac{\sqrt{\frac{p_f(y)}{p_{f^*}(y)}} e^{-\text{pen}(f)}}{\mathcal{A}(f^*, f)} \right] \right\} \\ & \leq \log \sum_{f \in \Gamma} \left[\frac{e^{-\text{pen}(f)}}{\mathcal{A}(f^*, f)} \underbrace{\mathbb{E}_{f^*} \left\{ \sqrt{\frac{p_f(y)}{p_{f^*}(y)}} \right\}}_{=\mathcal{A}(f^*, f)} \right] \\ & = \log \sum_{f \in \Gamma} e^{-\text{pen}(f)} \leq 0. \end{aligned}$$

By definition of \hat{f} , we have

$$\log \frac{p_{f^*}(y)}{p_{\hat{f}}(y)} + 2 \text{pen}(\hat{f}) \leq \min_{f \in \Gamma} \left[\log \frac{p_{f^*}(y)}{p_f(y)} + 2 \text{pen}(f) \right].$$

Thus

$$\begin{aligned} \mathbb{E}_{f^*} \left[\log \frac{p_{f^*}(y)}{p_{\hat{f}}(y)} + 2 \text{pen}(\hat{f}) \right] & \leq \min_{f \in \Gamma} \left[\mathbb{E}_{f^*} \log \frac{p_{f^*}(y)}{p_f(y)} \right. \\ & \quad \left. + 2 \text{pen}(f) \right] \\ & \equiv \min_{f \in \Gamma} \left[\text{KL}(p_{f^*} \| p_f) \right. \\ & \quad \left. + 2 \text{pen}(f) \right]. \end{aligned}$$

Putting everything together, we get the bound

$$2 \mathbb{E} \log \frac{1}{\mathcal{A}(f^*, f)} \leq \min_{f \in \Gamma} \left[\text{KL}(p(\cdot | A f^*) \| p(\cdot | A f)) + 2 \text{pen}(f) \right]$$

as advertised.

ACKNOWLEDGMENT

The authors would like to thank E. Candès for several fruitful discussions and the anonymous reviewers for very helpful suggestions.

REFERENCES

[1] E. J. Candès and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?," *IEEE Trans. Inf. Theory*, vol. 52, pp. 5406–5425, Dec. 2006.
 [2] D. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, pp. 1289–1306, Apr. 2006.
 [3] D. Lingenfelter, J. Fessler, and Z. He, "Sparsity regularization for image reconstruction with Poisson data," in *Proc. SPIE Computat. Imag. VII*, 2009, vol. 7246.

[4] Z. Harmany, R. Marcia, and R. Willett, "Sparse Poisson intensity reconstruction algorithms," in *Proc. IEEE Workshop on Statist. Signal Process.*, 2009 [Online]. Available: <http://arxiv.org/abs/0905.0483>
 [5] M. Figueiredo and J. Bioucas-Dias, "Deconvolution of Poissonian images using variable splitting and augmented Lagrangian optimization," in *Proc. IEEE Workshop on Statist. Signal Process.*, 2009 [Online]. Available: <http://arxiv.org/abs/0904.4868>
 [6] J.-L. Starck and J. Bobin, "Astronomical data analysis and sparsity: From wavelets to compressed sensing," *Proc. IEEE: Special Issue on Appl. Sparse Represent. Compress. Sensing*, 2010, to be published.
 [7] E. Candès, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Commun. Pure Appl. Math.*, vol. 59, no. 8, pp. 1207–1223, 2006.
 [8] J. Haupt and R. D. Nowak, "Signal reconstruction from noisy random projections," *IEEE Trans. Inf. Theory*, vol. 52, pp. 4036–4048, Sep. 2006.
 [9] D. L. Snyder, A. M. Hammond, and R. L. White, "Image recovery from data acquired with a charge-coupled-device camera," *J. Opt. Soc. Amer. A*, vol. 10, pp. 1014–1023, 1993.
 [10] R. Willett and R. Nowak, "Multiscale Poisson intensity and density estimation," *IEEE Trans. Inf. Theory*, vol. 53, pp. 3171–3187, Sep. 2007.
 [11] R. Willett and R. Nowak, "Platelets: A multiscale approach for recovering edges and surfaces in photon-limited medical imaging," *IEEE Trans. Med. Imag.*, vol. 22, pp. 332–350, 2003.
 [12] E. J. Candès and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 15, pp. 4203–4215, Dec. 2005.
 [13] J. A. Tropp, "Just relax: Convex programming methods for identifying sparse signals in noise," *IEEE Trans. Inf. Theory*, vol. 52, pp. 1030–1051, Mar. 2006.
 [14] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. Roy. Statist. Soc. Ser. B*, vol. 58, no. 1, pp. 267–288, 1996.
 [15] W. Dai and O. Milenkovic, "Weighted superimposed codes and constrained integer compressive sensing," *IEEE Trans. Inf. Theory*, vol. 55, pp. 2215–2229, May 2009.
 [16] M. A. Khajehnejad, A. G. Dimakis, W. Xu, and B. Hassibi, "Sparse recovery of positive signals with minimal expansion," 2009 [Online]. Available: <http://arxiv.org/abs/0902.4045>, submitted for publication
 [17] F. Bunea, A. Tsybakov, and M. Wegkamp, "Sparse density estimation with ℓ_1 penalties," in *COLT*, 2007.
 [18] I. Rish and G. Grabarnik, "Sparse signal recovery with exponential-family noise," in *Proc. Allerton Conf. Commun., Control, Comput.*, 2009.
 [19] S. Mendelson, A. Pajor, and N. Tomczak-Jaegermann, "Reconstruction and sub-Gaussian operators in asymptotic geometric analysis," *Geom. Funct. Anal.*, vol. 17, pp. 1248–1282, 2007.
 [20] S. Mendelson, A. Pajor, and N. Tomczak-Jaegermann, "Uniform uncertainty principle for Bernoulli and sub-Gaussian ensembles," *Construct. Approx.*, vol. 28, pp. 277–289, 2008.
 [21] M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, T. Sun, K. F. Kelly, and R. G. Baraniuk, "Single pixel imaging via compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 83–91, 2008.
 [22] M. Gehm, R. John, D. Brady, R. Willett, and T. Schultz, "Single-shot compressive spectral imaging with a dual-disperser architecture," *Opt. Expr.*, vol. 15, no. 21, pp. 14013–14027, 2007.
 [23] R. DeVerse, R. Coifman, A. Coppi, W. Fateley, F. Geshwind, R. Hamaker, S. Valenti, F. Warner, and G. Davis, "Application of spatial light modulators for new modalities in spectrometry and imaging," in *Proc. SPIE*, 2003, vol. 4959.
 [24] M. A. Neifeld and J. Ke, "Optical architectures for compressive imaging," *Appl. Opt.*, vol. 46, pp. 5293–5303, 2007.
 [25] S. Jafarpour, R. Willett, M. Raginsky, and R. Calderbank, "Performance bounds for expander-based compressed sensing with Poisson noise," in *Proc. Asilomar*, 2009.
 [26] K. Krishnamurthy, M. Raginsky, and R. Willett, "Multiscale photon-limited hyperspectral image reconstruction," *SIAM J. Imag. Sci.*, 2010, submitted for publication.
 [27] A. R. Barron, C. Huang, J. Q. Li, and X. Luo, "MDL, penalized likelihood, and statistical risk," in *Proc. IEEE Inf. Theory Workshop*, Porto, Portugal, May 2008, pp. 247–257.
 [28] I. Csiszár, "Why least squares and maximum entropy? An axiomatic approach to inference for linear inverse problems," *Ann. Statist.*, vol. 19, no. 4, pp. 2032–2066, 1991.
 [29] F.-X. Dupé, J. M. Fadili, and J.-L. Starck, "A proximal iteration for deconvolving Poisson noisy images using sparse representations," *IEEE Trans. Image Process.*, vol. 18, pp. 310–321, 2009.
 [30] S. Setzer, G. Steidl, and T. Teuber, "Restoration of images with rotated shapes," *Numer. Alg.*, vol. 48, no. 1–3, pp. 49–66, Jul. 2008.


- [31] P. L. Combettes and J.-C. Pesquet, "A proximal decomposition method for solving convex variational inverse problems," *Inverse Prob.*, vol. 24, no. 6, pp. 065014–065014, 2008.
- [32] C. Chau, J.-C. Pesquet, and N. Pustelnik, "Nested iterative algorithms for convex constrained image recovery problems," *SIAM J. Imag. Sci.*, vol. 2, no. 2, pp. 730–762, 2009.
- [33] A. W. van der Waat and J. A. Wellner, *Weak Convergence and Empirical Processes*. New York: Springer, 1996.
- [34] Q. Li and A. Barron, "Mixture density estimation," in *Advances in Neural Information Processing*. Cambridge: MIT Press, 2000, vol. 12.
- [35] E. Kolaczyk and R. Nowak, "Multiscale likelihood analysis and complexity penalized estimation," *Ann. Statist.*, vol. 32, pp. 500–527, 2004.
- [36] Y. Censor and S. A. Zenios, *Parallel Optimization: Theory, Algorithms, and Applications*. Oxford: Oxford Univ. Press, 1997.
- [37] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. New York: Wiley, 2006.
- [38] A. Bhattacharyya, "On a measure of divergence between two statistical populations defined by probability distributions," *Bull. Calcutta Math. Soc.*, vol. 35, pp. 99–109, 1943.
- [39] H. Chernoff, "A measure of asymptotic efficiency of test of a hypothesis based on the sum of observations," *Ann. Math. Statist.*, vol. 23, pp. 493–507, 1952.



Maxim Raginsky (S'99–M'00) received the B.S. and M.S. degrees in 2000 and the Ph.D. degree in 2002 from Northwestern University, Chicago, IL, all in electrical engineering.

From 2002 to 2004, he was a Postdoctoral Researcher with the Center for Photonic Communication and Computing, Northwestern University, where he pursued work on quantum cryptography and quantum communication and information theory. From 2004 to 2007, he was a Beckman Foundation Postdoctoral Fellow with the University

of Illinois in Urbana-Champaign, where he carried out research on information theory, statistical learning, and computational neuroscience. From 2007 to 2010, he has held the Research Scientist position with the Department of Electrical and Computer Engineering, Duke University, Durham, NC, where he is now an Assistant Research Professor. His general research interests pertain to information processing and decision-making in highly uncertain environments under resource and complexity constraints. His research combines techniques from information theory, statistical learning theory, control, optimization and games, and statistical signal processing.



Rebecca M. Willett (S'02–M'05) received the B.S.E. degree (*summa cum laude*) in electrical and computer engineering from Duke University, Durham, NC, in 2000, and the M.S. and Ph.D. degrees in electrical and computer engineering from Rice University, Houston, TX, in 2002 and 2005, respectively.

She has also held visiting researcher positions with the Institute for Pure and Applied Mathematics, University of California, Los Angeles (UCLA), in 2004, the University of Wisconsin-Madison 2003–2005, the French National Institute for Research in Computer Science and Control (INRIA) in 2003, and the Applied Science Research and Development Laboratory, GE Healthcare, in 2002. She is now an Assistant Professor of Electrical and Computer Engineering at Duke University. Her research interests include signal and image processing and communications with applications in medical imaging, astronomy, and wireless sensor networks.

Dr. Willett received the National Science Foundation CAREER Award in 2007.



Zachary T. Harmany (S'02) received the B.S. degree (*magna cum laude*) in electrical engineering and the B.S. degree (*cum laude*) in physics from The Pennsylvania State University, Philadelphia, in 2006.

Currently, he is a graduate student with the Department of Electrical and Computer Engineering, Duke University, Durham, NC. In 2010, he was a Visiting Researcher with The University of California, Merced. His research interests include nonlinear optimization, functional neuroimaging, and signal and image processing with applications in medical

imaging, astronomy, and night vision.



Roummel F. Marcia (M'08) received the B.A. degree in mathematics from Columbia University, New York, in 1995, and the Ph.D. degree from the University of California, San Diego, in 2002.

He was a Computational and Informatics in Biology and Medicine Postdoctoral Fellow with the Biochemistry Department, University of Wisconsin-Madison, and a Research Scientist with the Department of Electrical and Computer Engineering, Duke University, Durham, NC. He is currently an Assistant Professor of Applied Mathematics with

the University of California, Merced. His research interests include nonlinear optimization, numerical linear algebra, signal and image processing, and computational biology.