REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggesstions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any oenalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
19-09-2014	Final Report	1-Sep-2010 - 31-Aug-2014

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Final Report: High Information Capacity Quantum Imaging	W911NF-10-1-0404
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
	0D10BH

6. AUTHORS	5d. PROJECT NUMBER
John C. Howell	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES	8. PERFORMING ORGANIZATION REPORT NUMBER
University of Rochester 518 Hylan Bldg. Rochester, NY 14627 -0140	

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211	ARO
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)
	58494-PH-DRP.13

12. DISTRIBUTION AVAILIBILITY STATEMENT

Approved for Public Release; Distribution Unlimited

13. SUPPLEMENTARY NOTES

The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT

This is the final report for the DARPA InPho program. In reality, we finished this program in early 2013, but for some reason, the program stayed open until August 2014 through the Army. We filed a final report for the 2013 interim report. Please see that report.

15. SUBJECT TERMS

Quantum Imaging, Photon Counting, LIDAR

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			John Howell
UU	UU	UU	UU		19b. TELEPHONE NUMBER
					585-275-0538

Report Title

Final Report: High Information Capacity Quantum Imaging

ABSTRACT

This is the final report for the DARPA InPho program. In reality, we finished this program in early 2013, but for some reason, the program stayed open until August 2014 through the Army. We filed a final report for the 2013 interim report. Please see that report.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

<u>Received</u>   <u>Paper</u>

08/28/2013 4.00 James Schneeloch, P. Ben Dixon, Gregory A. Howland, Curtis J. Broadbent, John C. Howell. Violation of Continuous-Variable Einstein-Podolsky-Rosen Steering with Discrete Measurements,
Physical Review Letters, (03 2013): 0. doi: 10.1103/PhysRevLett.110.130407

08/28/2013 2.00 Gregory A. Howland, John C. Howell. Efficient High-Dimensional Entanglement Imaging with a Compressive-Sensing Double-Pixel Camera,
Physical Review X, (02 2013): 0. doi: 10.1103/PhysRevX.3.011013

08/28/2013 3.00 Gregory A. Howland, James Schneeloch, John C. Howell, P. Ben Dixon. Quantum Mutual Information Capacity for High-Dimensional Entangled States,
Physical Review Letters, (04 2012): 0. doi: 10.1103/PhysRevLett.108.143603

08/28/2013 5.00 James Schneeloch, Curtis J. Broadbent, Stephen P. Walborn, Eric G. Cavalcanti, John C. Howell. Einstein-Podolsky-Rosen steering inequalities from entropic uncertainty relations,
Physical Review A, (06 2013): 0. doi: 10.1103/PhysRevA.87.062103

08/28/2013 6.00 Omar S. Magan?a-Loaiza, Gregory A. Howland, Mehul Malik, John C. Howell, Robert W. Boyd. Compressive object tracking using entangled photons,
Applied Physics Letters, (06 2013): 0. doi: 10.1063/1.4809836

08/28/2013 7.00 Eric Lantz, Jeffrey H. Shapiro. Comment on "Observation of anticorrelation in incoherent thermal light fields",
Physical Review A, (05 2012): 0. doi: 10.1103/PhysRevA.85.057801

08/28/2013 8.00 Kam Wai Clifford Chan, D. S. Simon, A. V. Sergienko, Nicholas D. Hardy, Jeffrey H. Shapiro, P. Ben Dixon, Gregory A. Howland, John C. Howell, Joseph H. Eberly, Malcolm N. O'Sullivan, Brandon Rodenburg, Robert W. Boyd. Theoretical analysis of quantum ghost imaging through turbulence,
Physical Review A, (10 2011): 0. doi: 10.1103/PhysRevA.84.043807

08/28/2013 9.00 Nicholas D. Hardy, Jeffrey H. Shapiro. Computational ghost imaging versus imaging laser radar for three-dimensional imaging,
Physical Review A, (02 2013): 0. doi: 10.1103/PhysRevA.87.023820

08/28/2013 10.00 Zhimin Shi, Malcolm N. O'Sullivan, Petros Zerom, Kam Wai Clifford Chan, Molly Krogstad, Jeffrey H. Shapiro, Robert W. Boyd. Thermal ghost imaging with averaged speckle patterns,
Physical Review A, (12 2012): 0. doi: 10.1103/PhysRevA.86.063817

08/28/2013 11.00 G. A. Howland, P. B. Dixon, J. C. Howell. Photon-counting compressive sensing laser radar for 3D imaging,
Applied Optics, (10 2011): 0. doi: 10.1364/AO.50.005917

08/28/2013 12.00 P. Ben Dixon, Gregory A. Howland, Kam Wai Clifford Chan, Colin O'Sullivan-Hale, Brandon Rodenburg, Nicholas D. Hardy, Jeffrey H. Shapiro, D. S. Simon, A. V. Sergienko, R. W. Boyd, John C. Howell. Quantum ghost imaging through turbulence,
Physical Review A, (05 2011): 0. doi: 10.1103/PhysRevA.83.051803

TOTAL: **11**

Number of Papers published in peer-reviewed journals: 4

(b) Papers published in non-peer-reviewed journals (N/A for none)

<u>Received</u>   <u>Paper</u>

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u>   <u>Paper</u>

TOTAL:

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u>   <u>Paper</u>

TOTAL:

(d) Manuscripts

<u>Received</u>   <u>Paper</u>

   TOTAL:

Number of Manuscripts:

Books

<u>Received</u>   <u>Book</u>

   TOTAL:

<u>Received</u>   <u>Book Chapter</u>

   TOTAL:

Patents Submitted

Patents Awarded

Awards

Graduate Students

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Names of Post Doctorates

NAME	PERCENT_SUPPORTED
James Schneel	0.39
FTE Equivalent:	**0.39**
Total Number:	**1**

Names of Faculty Supported

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Names of Under Graduate students supported

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Student Metrics
This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 0.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:...... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:...... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):...... 0.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:...... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields: 0.00

Names of Personnel receiving masters degrees

NAME
Total Number:

Names of personnel receiving PHDs

NAME
Gregory Howland
Total Number: **1**

Names of other research staff

NAME   PERCENT_SUPPORTED

FTE Equivalent:
Total Number:

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

Please see the final report that was filed as last years interim report.

Technology Transfer

Final Report on DARPA InPho High Capacity Quantum Imaging

Robert W. Boyd, John C. Howell

Department of Physics and Astronomy, University of Rochester, Rochester, New York, 14627, USA

Vivek Goyal, Franco N. C. Wong, and Jeffrey H. Shapiro

Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139

Alexander V. Sergienko

Dept. of Electrical and Computer Engineering,
Boston University, 8 Saint Mary's St., Boston, MA 02215

(Dated: August 30, 2013)

We report on highlights within the DARPA InPho program, which come closest to meeting the end-of-program goals. First, we demonstrate a novel parametric process imaging technique with greater than 1 bit per photon by reconstructing a high resolution depth map of a natural scene. Second, we then discuss a double-pixel compressive sensing architecture for characterizing high-dimensional transverse entanglement showing greater than 8 bits per photon pair. Third, we demonstrate a high resolution linear compressive sensing depth map, which achieved high frame rate object tracking in 3 dimensions. Fourth, we report on ghost imaging techniques and point target localization at greater at 1 bit per photon. Lastly, we discuss the use of orbital angular momentum decomposition imaging.

CONTENTS

I. PARAMETRIC POISSON PROCESS IMAGING

We introduce a new conceptual framework for active optical imaging that leads to striking capabilities in low-light scenarios. We consider the use of raster-scanned illumination and detection of light with a single-photon avalanche diode (SPAD); the framework applies identically with flood illumination and a SPAD array. Due to the quantum nature of light, the SPAD measures a Poisson process at each spatial location. The framework is called *parametric Poisson process imaging* (PPPI) because the essential idea is that the rate of each Poisson process is a parametric function with parameters that encode features such as scene reflectance, distance to scene points, and background light level. The parametric nature of these functions along with correlations over the spatial dimensions enable effective estimation with only a few detected photons and with significant contribution from background light.

A. Main contributions

Conceptual: We introduce a physically-accurate model for the signal produced by a SPAD under weak lighting which incorporates arbitrary illumination pulse shape, inhomogeneous Poisson process characteristics (shot noise from the quantum nature of light), and detector reset time.

Algorithmic: We provide a method for spatiotemporally-regularized estimation of intensity and depth. Our algorithm produces accurate intensity and depth images from 1 photon per pixel (ppp) data.

Experimental: In work detailed elsewhere [1], we simultaneously achieve very high photon efficiency for imaging intensity and depth using raster-scanned illumination with a pulsed laser and detection of the backscattered light with a single SPAD.

- Intensity estimation for a 16-level scene at 1 ppp is comparable to a traditional technique with 200 ppp.

- Depth estimation at 1 ppp is comparable to a traditional technique with 100 ppp or 17 times narrower pulse width.

B. Related Work

Algorithmic: Image processing at low-light levels is challenging due to the signal dependent nature of noise. Denoising of low-light images captured by SPAD sensors is typically accomplished using the Poisson or shot noise model along with sparsity promoting regularization [2]. These methods assume that the number of collected photons is Poisson distributed with parameter proportional to the true intensity. This signal model fails when the photon count is too low, such as in our case. For depth imaging using SPAD data, first a pixelwise maximum likelihood (ML) estimate of scene depth is computed using a time-inhomogeneous Poisson process model for photon arrival times [3, 4]. Then sparsity-promoting regularization is used to further denoise this depth map. This two-step approach assumes a Gaussian noise model which is befitting because of the behavior of ML with large number of data samples [5]. The regime of image processing using very low photon counts, which is the focus of this work, is largely unexplored. Another challenging algorithmic problem that is not addressed in the existing literature is the filtering of stray photons due to ambient light and detector dark counts.

Depth imaging using active illumination: Several techniques differ in the way active illumination is modulated. Time-intensity modulated imagers operate on the TOF principle. When ordered by increasing modulation bandwidth (shorter pulses), these are: homodyne TOF sensing [6], pulsed TOF cameras [7], and picosecond laser radar systems [8]. Sensors that spatially modulate the light include speckle decorrelation imaging, structured light, and active stereo imaging [9]. Intensity modulated sensing requires specialized detectors, whereas spatially-modulated sensing is accomplished using standard CMOS detector arrays. However, spatially-modulated detectors have low photon efficiency because they achieve centimeter range resolution using high optical output (always on). In comparison, TOF systems achieve millimeter accurate sensing by using lower optical output (on for a short time). Among the aforementioned systems, TOF imagers using
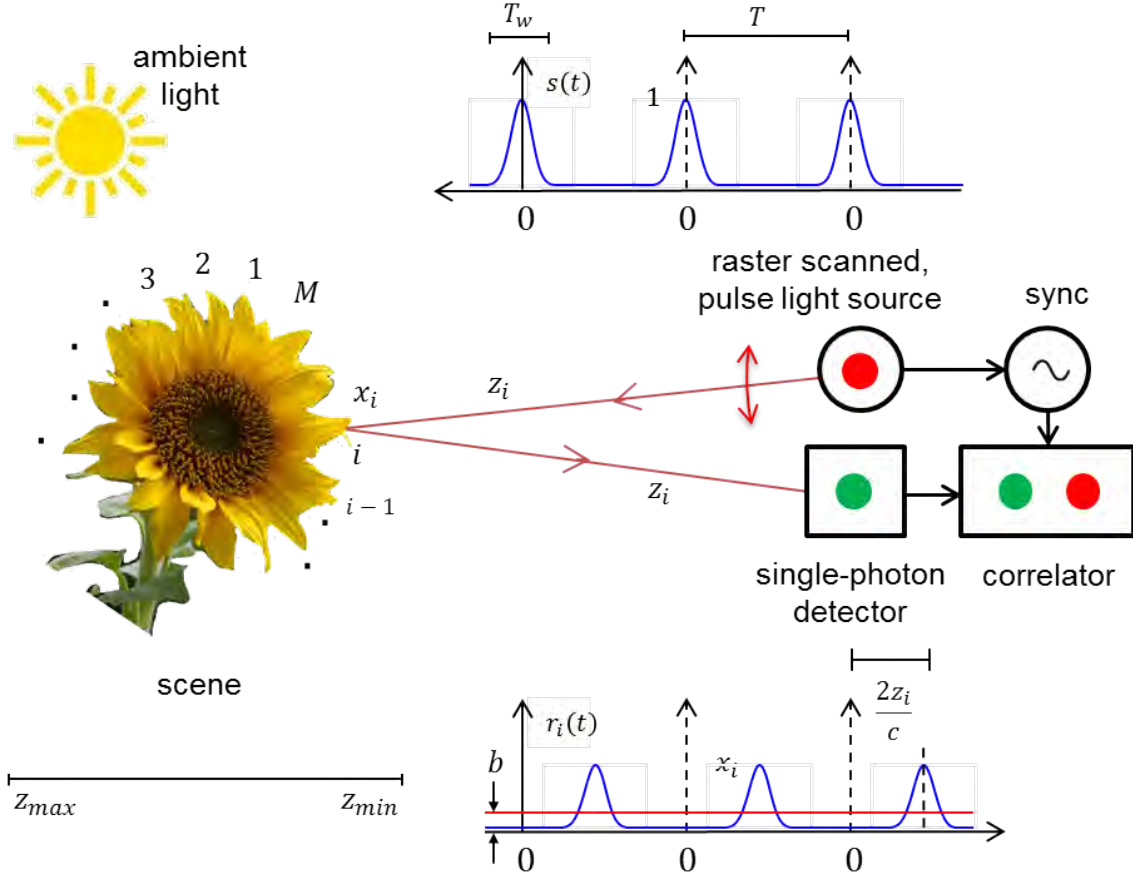
FIG. 1. Our imaging setup.

SPAD detectors have the highest photon efficiency. Our proposed framework further improves the high photon efficiency of SPAD imagers, which translates to lower optical power and lower system bandwidth without sacrificing imaging quality.

C. Imaging Setup

Our imaging setup is shown in Figure 1. The light source and SPAD are co-located. Scene patches are indexed by $i \in \{1, 2, \ldots, M\}$. The reflectance of patch i, denoted x_i, includes the effect of radial fall-off, viewing angle and material properties. The distance to patch i is denoted z_i.

We use an intensity-modulated light source with pulse shape $s(t)$ and repetition interval T. We assume that $T > 2 z_{\max}/c$, where $z_{\max}$ is largest scene depth and the c is the speed of light; this avoids distance aliasing. With conventional processing, the pulse width T_w governs the achievable depth resolution [3]. We assume $T_w \ll 2 z_{\max}/c < T$. Denote the total intensity of a single pulse by $S = \int_0^T s(t)\, dt$.

Light incident on the detector is spectrally filtered to the wavelength of the source so that ambient light is largely rejected. The fraction of photons passing through the filter that trigger electrical impulses is the detector quantum efficiency η. Every detected photon is time stamped within an accuracy of Δ, called the *time bin size*. We assume that $\Delta \ll T_w \ll 2 z_{max}/c$. When the detector records a photon arrival—referred to as a *click*—it becomes inactive for a small duration called the reset time or dead time. As long as the reset time is longer than T, it does not affect our modeling.
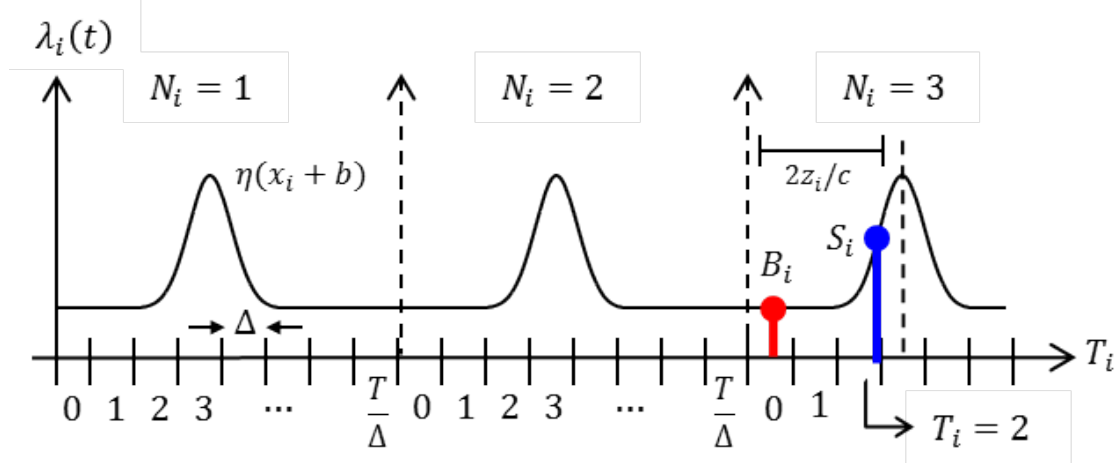
FIG. 2. Rate function of inhomogeneous Poisson process combining desired scene response and noise sources. The detected signal photon (blue) was generated by the third transmitted pulse ($N_i = 3$) and arrived in the second time slot ($T_i = 2$). An arrival from the noise process is shown in red.

For each raster position i, we repeatedly illuminate scene patch i until the first click. The discrete time bin T_i of the click (relative to the pulse that resulted in the click) and the number of pulses up to and including the click N_i are retained, so the data set is $\{(T_i, N_i)\}_{i=1}^{M}$.

D. Measurement Model

Illuminating scene patch i with intensity-modulated light pulse $s(t)$ results in a backscattered light signal $r_i(t)$ that is physically modeled as $r_i(t) = x_i\, s(t - 2z_i/c) + b_\lambda$, where b_λ denotes the background light intensity at the operating wavelength. The function $\eta\, r_i(t)$ is the time-varying rate of an inhomogeneous Poisson process observed at the SPAD, and photon arrivals are recorded per these statistics. The SPAD detector records these photon arrivals in time bins of size Δ. It is also necessary to add the detector dark counts, which are modeled as an independent time-invariant Poisson process with rate d. The observed inhomogeneous Poisson process thus has rate

$$\lambda_i(t) = \eta\, r_i(t) + d = \eta\, x_i\, s(t - 2z_i/c) + (\eta b_\lambda + d).$$

By defining $b = b_\lambda + d/\eta$, we see that η is an overall multiplicative factor that can be set to 1 without loss of generality. We assume that b is known since it is straightforward to estimate. Repetitions until the first click create a periodic rate function as shown in Figure 2.

Low-rate assumption: Our computations of the distributions of random variables T_i and N_i hold when the light intensity at the detector is low, as would be the case when photon efficiency is important. Mathematically, the number of detected photons from a single pulse (neglecting detector dead time) is a random variable; let A be the event that this variable is at most 1. Conditioning on A is important because it allows the decoupling of intensity and depth estimation.

Time bin T_i: The depth of patch i is encoded in the bin index T_i. It is straightforward to derive the conditional probability mass function (PMF) of T_i given A and approximate it as

$$p_{T_i|A}(k_i) \;\approx\; \lambda(k_i\Delta)\Delta\Lambda,$$

where $\Lambda = \int_0^T \lambda(t)\,dt = x_i S + bT$, for small Δ. This has an intuitive interpretation: the probability of the photon being detected in time bin k_i is proportional to the area of the rate function which falls within the time bin.
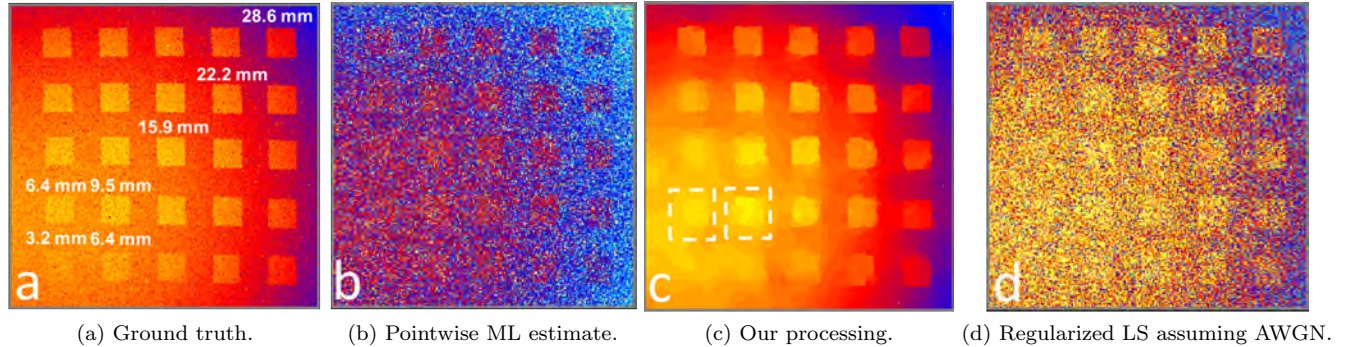
(a) Ground truth.　　　(b) Pointwise ML estimate.　　　(c) Our processing.　　　(d) Regularized LS assuming AWGN.

FIG. 3. **Depth resolution test from 1 ppp data.** Thin squares in the lower-left corner of (a)/(c) are completely missing from (b)/(d).

(a) Ground truth.　　　(b) Pointwise ML estimate.　　　(c) Our processing.　　　(d) Regularized LS assuming AWGN.
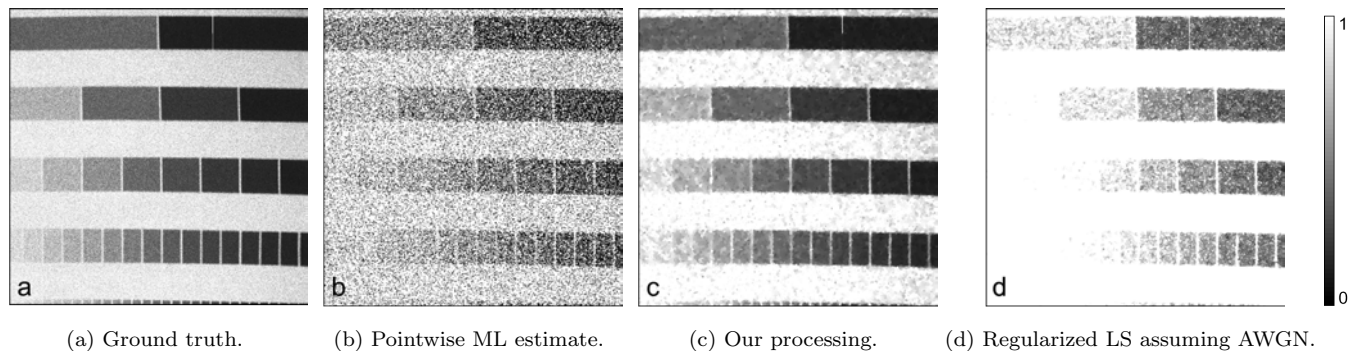
FIG. 4. **Intensity resolution test from 1 ppp data.** More gray levels (14) are distinguishable in (c) than in (d).

Pulse count N_i: Novel to PPPI, we exploit that the intensity of patch i is encoded in the number of pulses transmitted before the first click N_i. The conditional PMF of N_i is

$$p_{N_i|A}(n_i) \ = \ e^{-\Lambda(N_i-1)}(1-e^{-\Lambda}), \qquad n_i = 1, \, 2, \, \ldots, \tag{1}$$

a geometric distribution with parameter $e^{-\Lambda} = e^{-(x_i S + bT)}$.

Signal vs. noise: A detected photon could have originated from the scene, or it could be noise due to ambient light or detector dark counts. For any pixel index i, let the event S_i be the detection of a signal photon (as opposed to noise). We use $\mathcal{S}$ to denote the set of pixel indices which have signal photons, and PPPI includes a step to estimate $\mathcal{S}$.

E.　Novel Image Formation

Our proposed image formation methods are regularized ML estimators. We combine physically-accurate probabilistic signal models derived in the previous section with regularization using sparsity. With this approach, we achieve significant improvements over pointwise ML and over traditional denoising as post-processing applied without a detailed physical model.

Intensity estimation: Using (1), we define the negative log-likelihood function for estimating scene reflectance x_i from data $N_i = n_i$ as

$$\mathcal{L}_x(x_i; n_i) = S(N_i - 1)\, x_i - \log(1 - e^{-(x_i S + bT)}).$$

This is a strictly convex function in x_i, making it amenable for global minimization using convex optimization. The regularized ML estimate for scene reflectance is obtained from noisy data $\{N_i = n_i\}_{i=1}^M$ by solving the
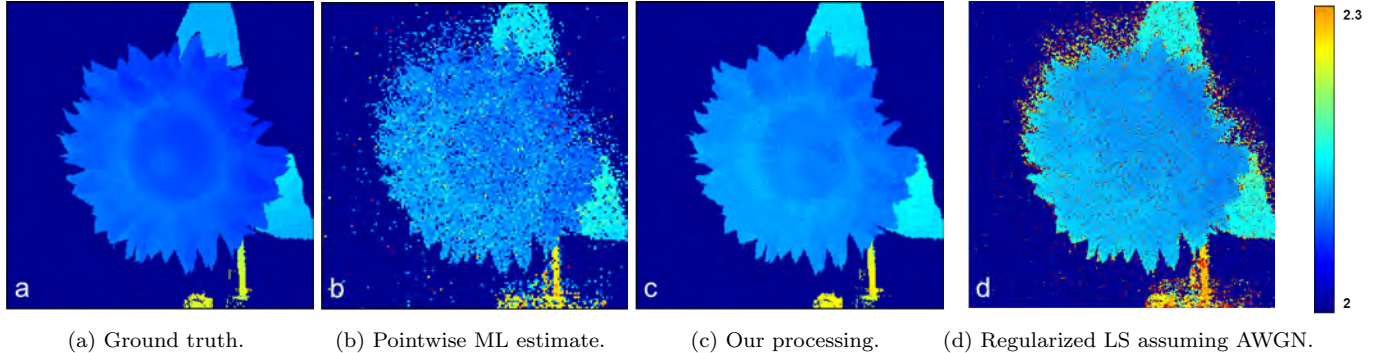
(a) Ground truth.   (b) Pointwise ML estimate.   (c) Our processing.   (d) Regularized LS assuming AWGN.

FIG. 5. **Sunflower depth map.** Our method rejects background and denoises while retaining fine spatial features like flower petals.

(a) Ground truth.   (b) Pointwise ML estimate.   (c) Our processing.   (d) Regularized LS assuming AWGN.
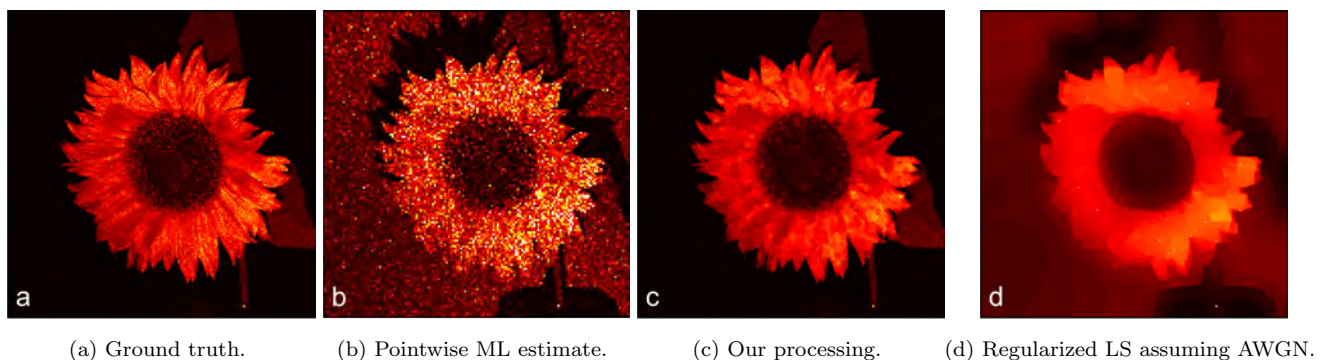
FIG. 6. **Sunflower intensity image.** Our method increases image contrast while retaining fine spatial features like flower petals.

following convex problem:

$$\min_{\substack{X=[x_1,\ldots,x_M] \\ 0 \leq x_i \leq 1}} \sum_{i=1\ldots M} \mathcal{L}_x(x_i; n_i) + \beta \left\| \Phi_x X \right\|, \tag{2}$$

where Φ_x is a sparsifying transform and different norms may be used.

Noise rejection: Noise events are generated from a homogeneous Poisson process with constant rate, resulting in a uniform distribution for T_i. In contrast, signal photons are generated by a time-inhomogeneous Poisson process whose rate is proportional to the backscattered waveform (or time duration T_w). Thus, signal photon arrivals result in a distribution with much lower variance compared with the uniform distribution over $[0, T]$. In fact, it is effective to treat the conditional distribution of T_i given S_i as impulsive.

It is shown in [10] that the rank-ordered absolute differences (ROAD) statistic gives a good characterization of the impulse distribution. Using the algorithm from [10], we filter noise using the ROAD statistic at each pixel i computed from its 8 nearest neighbors. The remaining signal photons form set $\mathcal{S}$.

Depth estimation: Once noise is rejected, we can readily formulate the negative log-likelihood function for estimation depth at pixel $i \in \mathcal{S}$ from data $T_i = k_i$ as

$$\mathcal{L}_z(z_i; k_i) = -\log \Pr[T_i = k_i | A, S_i] \equiv -\log s(k_i \Delta - 2z_i/c).$$

Our framework allows the use of arbitrary pulse shapes, but many practical pulse shapes like Gaussian, and lopsided pulse shapes are well approximated as $s(t) = e^{-f(t)}$, where $f(t)$ is a convex function in t. In such cases, $\mathcal{L}_z(z_i; k_i) = f(k_i \Delta - 2z_i/c)$ is a convex function in z_i. Regardless of convexity of $f(t)$, global optimization is still possible for arbitrary pulse shapes if a good initial starting point is available. In our

case, the signal photons' arrival times serve as excellent starting points because of their proximity to the true depth estimate. Starting with these initial values, we use gradients of the pulse shape $s(t)$ along with sparsity-based regularization to obtain close to globally-optimal solutions. Our regularized ML estimate for scene depth is obtained from filtered data $\{T_i = k_i\}_{i \in S}$ by solving the following optimization problem:

$$\begin{array}{c} Z=[z_1,\ldots,z_M] \\ z_{\min} \leq z_i \leq z_{\max} \end{array} \sum_{i \in S} \mathcal{L}_z(z_i; k_i) + \beta \left\| \Phi_z Z \right\|_\ell. \tag{3}$$

F.   Experimental Results

We conducted experiments, to demonstrate intensity and depth estimation accuracy of our framework; experimental parameters are detailed in [1]. For reflectance estimation, we used ℓ_1 regularization with 4-tap Daubechies wavelet basis. For depth estimation, we used TV regularization. The regularization parameter β for each method was chosen to maximize PSNR for the given scene.

As shown in Figure 3, our method achieves a depth resolution between 4 mm and 6 mm using 1 ppp. To put our results in context, the pulse width required to achieve 5 mm depth resolution in a noiseless setting is ≈ 17 ps; in our experiments we used noisy data and a pulse width of 1 ns, which is 60 times longer. As expected, regularized AWGN ML estimation fails due to impulsive noise, and aggressive regularization leads to oversmoothing and loss of spatial features. As indicated by our results in Figure 4, our proposed method resolves about 14 gray levels using 1 ppp data, far better than the raw data or regularized LS under an AWGN model.

The estimation results for a natural scene are shown in Figures 5 and 6. As shown, our framework achieves a significant gain in depth resolution and boost in PSNR relative to conventional methods while preserving high frequency details and small features.

II.   COMPRESSIVE HIGH-DIMENSIONAL ENTANGLEMENT CHARACTERIZATION

Spatially entangled biphotons, such as those generated by spontaneous parametric downconversion (SPDC), exhibit strong Einstein-Podolsky-Rosen (EPR) type correlations [11] in the transverse position and transverse momentum degrees of freedom [12]. Because these variables are continuous, the entanglement can be very high-dimensional with a typical Schmidt number greatly exceeding 1000 [13]. This provides high information density which can be leveraged to increase channel capacity and security for quantum key distribution [14–16] and dense coding [17, 18]. Other applications include ghost imaging [19, 20], quantum computing [21], and quantum teleportation [22].

Experimentally characterizing the SPDC state is unfortunately difficult due to weak sources and low resolution detectors. Spatial entanglement is traditionally imaged by jointly raster-scanning photon-counting avalanche photodiodes (APDs) to measure spatial correlations. This scales extremely poorly with increasing detector resolution. With a biphoton flux of $4,000$ coincident detections per second, it would take 55 days to jointly scan a 24×24 pixel region for a signal-to-noise ratio (SNR) of 10. For 32×32 pixels, it would take 310 days (see Eq. 12).

Other approaches have been tried with mixed success. Intensified CCD cameras can measure the Schmidt number [23], but do not detect single-photon correlations rendering them ineffective for most quantum applications. Arrays of photon counting APDs could replace CCDs, but they are currently low resolution, noisy and resource intensive; especially since each pixel pair must be individually correlated [24–26]. A recent, promising result averages intensity correlations over many images from a single-photon sensitive electron-multiplying CCD, reporting 2500 modes [27]. This technique is limited to a 30 ms exposure time (APDs are sub-ns) and is noisier than using APDs because it does not isolate individual coincident detections.

In ref [28], Dixon *et. al.* reduce the number of measurements required for a raster-scan by only measuring in an area of interest where correlations are expected, reporting a channel capacity of 7 bits per photon. While not a true full-field measurement, they highlight a critical feature of the SPDC field. In *both* position and momentum representations, the distribution of correlations between pairs of detector pixels is very sparse despite dense (not sparse) single-particle distributions. Applying ideas from the field of compressive

sensing, we exploit prior knowledge of this sparsity to beat the "curse of dimensionality" [29] and efficiently characterize the full biphoton field without raster scanning.

Here, we implement a compressive sensing, photon-counting double-pixel camera that efficiently images single-photon SPDC correlations in the near- and far-field at resolutions up to $32 \times 32 = 1024$ dimensions per detector. At 32×32 resolution, the measurement time is reduced from 310 days for raster scanning to around 8 hours. We perform an entropic characterization showing channel capacities of up to 8.4 bits per photon, equivalent to 337 independent, identically distributed modes. Sums of channel capacities in conjugate bases violate a EPR steering bound [30] by up to 6.6 bits.

A.   Linear Compressive Sensing Theory

Compressive sensing is a technique that employs optimization to measure a sparsely represented N-dimensional signal from $M < N$ incoherent measurements [33, 34, 81, 82]. The approach is so-named because the signal is effectively compressed during measurement. Though sparsity is assumed, it is not known prior to measurement which elements contain appreciable amplitude. Compressive sensing must determine both which elements are significant and find their values.

To detect a sparsely represented N-dimensional signal vector X, we measure a series of $M < N$ values Y by multiplying X by an $M \times N$ sensing matrix A such that

$$Y = AX + \Gamma, \tag{4}$$

where Γ is a noise vector.

Because $M < N$, this system is undetermined; a given Y does not specify a unique X. The correct X is recovered by minimizing a regularized least squares objective function

$$\min_X \frac{1}{2} ||Y - AX||_2^2 + \tau g(X), \tag{5}$$

where for example $||\Omega||_2^2$ is the ℓ_2 norm of Ω and τ is a scaling constant. The function $g(X)$ is a regularization promoting sparsity. Common $g(X)$ include X's ℓ_1 norm, assuming the signal is sparse, and X's total variation, assuming the signal's gradient is sparse [35]. A must be incoherent with the basis of interest, with the surprising and non-intuitive result that a random, binary sensing matrix works well. Given sufficiently large M, the recovered X approaches the exact signal with high probability [36]. For a $k-$sparse signal, the required M scales as $M \propto k \log(N/k)$.

Incoherent, random sampling is particularly beneficial for low-light measurements as each measurement receives on average half the total photon flux $\Phi/2$, as apposed to Φ/n for a raster scan. Compressive sensing is now beginning to be used for quantum applications such as state tomography [37]. Shabani et. al, for example, perform a tomography on a two qubit photonic gate for polarization entangled photons [38]. CS has also been used to with spatially correlated light for ghost imaging [39, 40]. It is important to note that for ghost imaging, CS is not required to recover the full two-particle probability distribution as in entanglement characterization.

The quintessential compressed sensing example is the single-pixel camera [41, 75]. An object is imaged onto a Digital Micromirror device (DMD), a 2D binary array of individually-addressable mirrors that reflect light either to a single detector or a dump. Rows of the sensing matrix A consist of random, binary patterns placed sequentially on the DMD. For an N-dimensional image, minimizing Eq. 5 recovers images using as few as $M = 0.02N$ measurements.

B.   Compressive Sensing for Measuring Correlations

The single-pixel camera concept naturally adapts to imaging correlations by adding a second detector. Consider placing separate DMDs in the near-field or far-field of the SPDC signal and idler modes, where

"on" pixels are redirected to photon counting modules. The signal of interest is

$$p_x(u, v) = \int_u \mathrm{d}\vec{x_s} \int_v \mathrm{d}\vec{x_i} |\psi(\vec{x_s}, \vec{x_i})|^2 \tag{6}$$

$$p_k(u, v) = \int_u \mathrm{d}\vec{k_s} \int_v \mathrm{d}\vec{k_i} |\tilde{\psi}(\vec{k_s}, \vec{k_i})|^2, \tag{7}$$

where $p(u, v)$ represents the probability of a coincident detection between the u^{th} pixel on the signal DMD and v^{th} pixel on the idler DMD. The functions $\psi(\vec{x_s}, \vec{x_i})$ and $\tilde{\psi}(\vec{k_s}, \vec{k_i})$ are approximate position and momentum wavefunctions for the biphoton

$$\psi(\vec{x_s}, \vec{x_i}) = \mathcal{N} \exp\left(-\frac{(\vec{x_s} - \vec{x_i})^2}{4\sigma_c^2}\right) \exp\left(-\frac{(\vec{x_s} + \vec{x_i})^2}{16\sigma_p^2}\right)$$

$$\tilde{\psi}(\vec{k_s}, \vec{k_i}) = (4\sigma_p\sigma_c)^2 \mathcal{N} \exp(-\sigma_c^2(\vec{k_s} - \vec{k_i})^2)$$
$$\times \exp(-4\sigma_p^2(\vec{k_s} + \vec{k_i})^2). \tag{8}$$

Subscripts s and i refer to signal and idler photons respectively, σ_p and σ_c are the pump and correlation widths, and $\mathcal{N}$ is a normalizing constant. X of Eq. 5 is simply a one-dimensional reshaping of p_x or p_k.

Like the single-pixel camera, a series of random patterns are placed on the DMDs to form rows of A. For each pair of patterns, correlations between the signal and idler photons form the measurement vector Y. Minimization of Eq. 5 recovers $p(u, v)$.

While a fully random A is preferred, the DMDs only act on their respective signal or idler subspace, which prevents arbitrary A. Rows of A are therefore outer products of rows of single-particle sensing sensing matrices $\mathbf{a}$ and $\mathbf{b}$

$$A = \begin{pmatrix} \mathbf{a}_1 \otimes \mathbf{b}_1 \\ \mathbf{a}_2 \otimes \mathbf{b}_2 \\ \vdots \\ \mathbf{a}_m \otimes \mathbf{b}_m \end{pmatrix}, \tag{9}$$

where rows of $\mathbf{a}$ represent random patterns placed on the signal DMD and rows of $\mathbf{b}$ represent random patterns placed on the idler DMD. To make signal and idler photons distinguishable, $\mathbf{a}$ and $\mathbf{b}$ are not the same. The validity of Kronecker-type sensing matrices has been established and is of current interest in the CS community as attention shifts to higher dimensional signals [43, 44]. The measurement vector Y is obtained by counting coincident detections for the series of DMD configurations given by A.

A variety of reconstruction algorithms exist for Eq. 5, with their computational complexity dominated by repeatedly calculating AX and $A^T Y$ [45]. This is especially unwieldy for correlation measurements as the size of A is $M \times n^2$ for $n-$pixel DMDs. Using properties for Kronecker products [46], these can be more efficiently computed by

$$AX = \mathbf{diag}(\mathbf{b}\,\mathbf{sq}(X)\,\mathbf{a}^T) \tag{10}$$

$$A^T Y = \mathbf{vec}(\mathbf{b}^T\,\mathbf{od}\,(Y)\mathbf{a}), \tag{11}$$

where $\mathbf{sq}$ and $\mathbf{vec}$ reshape a vector to a square matrix and vice-versa; $\mathbf{diag}$ forms a vector from the diagonal elements of a square matrix; and $\mathbf{od}$ forms a square matrix placing the operand vector on its diagonal.

C.   Comparison to Raster Scanning

The compressive approach finds the joint probability distribution orders of magnitude faster than raster scanning through two key improvements. The first is simply the reduction in the number of measurements.

To jointly raster scan an n-pixel space requires n^2 measurements. For a compressive measurement, sparsity is approximately n with dimensionality n^2, so only $M \propto n \log(n)$ measurements are required. In practice, we found excellent results when M was only three percent of n^2.

The second advantage of compressive measurements is that they more efficiently use available flux. For the raster scan, the total flux is distributed over at best n pairs of pixels in the case of perfect correlations. Conversely, the average flux per incoherent compressive measurement is independent of n, with each measurement receiving on average $1/4$ the total flux. To maintain constant SNR (photons/measurement) with increasing n, total measurement time therefore scales as n^3 for raster scanning. Given a photon flux of Φ photons per second, the measurement time for a desired SNR is

$$t = n^2 t_{\text{meas}} = \frac{n^3 \text{SNR}^2}{\Phi}, \tag{12}$$

where t_{meas} is the time per measurement.

For incoherent, compressive measurements, acquisition time scales as $n \log(n)$. The compressive improvement therefore scales as $n^2 / \log(n)$. For $n = 32 \times 32 = 1024$, this is of order 10^5.

This scaling factor somewhat optimistically assumes the reconstruction process yields an accurate result despite a noisy signal. Unfortunately, propagation of uncertainty through the reconstruction process remains a difficult problem, especially for non-ideal, real world systems [47]. There has been much recent theoretical work on the topic for Gaussian [48–50] and Poissonian noise [51, 52]. These results tend to require ideal sensing matrices or more complicated formulations to give provable performance bounds. As such, their findings are difficult to directly, quantitatively apply to experiment. However, they do reveal pertinent features that indicate CS can perform extremely well in the presence of noise.

A well known characteristic of CS is a rapid phase change from poor to good quality reconstructions [53]. This phase change is often discussed as a function of increasing m, with the boundary $m \propto k \log(n/k)$. A similar phase transition occurs as a function of the noise level; in our case, this is the number of photons per measurement. For some cases, these two are even linked [48]. A practical compressive measurement simply requires large enough m and photon flux Φ to be in the space of good reconstructions. Fortunately, simply obtaining a recognizable reconstruction generally indicates the measurement conditions exceed this threshold.

Unlike a direct measurement, the information obtained by a series of y compressive measurements is contained in their deviation from the average value $\bar{y}$. In the presence of noise, these deviations must exceed the noise level. Assuming Poissonian shot noise, good reconstructions require $\text{std}(y) \geq \beta \sqrt{\bar{y}}$, where $\text{std}(y)$ is the standard deviation of the measurement vector and β is a positive constant greater than one.

The particular algorithm chosen to solve Eq. 5 also plays a role in the reconstruction's accuracy. These often have provable performance on ideal signals, but degrade when confronted with noisy or otherwise non-ideal conditions. In these circumstances, they have various strengths, including speed, accuracy, and sensitivity to user selected parameters such as τ in Eq. 5. For more information on common reconstruction algorithms, see Refs [34, 45, 55, 87].

In practice, the best way to determine accuracy for a particular signal, sensing matrix, and reconstruction approach remains repeated simulations or experiments. For our system, we reduce a $n = 32 \times 32$ measurement from a 310 day raster scan (SNR of 10) to an 8 hour compressive acquisition, a thousand-fold improvement.

D. Experiment

The experimental apparatus is given in Fig. 14. Light from a 2.8 mW, 325 nm HeCd laser was directed to a 1 mm long BiBO crystal oriented for type I, collinear SPDC. The generated daughter photons passed through a 650/13 nm narrowband filter before separating into signal and idler modes at a 50/50 beamsplitter. To measure position-position correlations, lenses $f1 = 125$ mm and $f2 = 500$ mm imaged the crystal onto signal and idler mode DMDs. For momentum-momentum correlations, $f1$ was removed and the DMDs were placed in the focal plane of $f2 = 88.3$ mm. DMD "on" pixels reflected light to large area, single photon counting modules (SPCM) connected to a correlating circuit.

To measure $p(u, v)$, a series of M random patterns were placed on the DMDs to form the sensing matrix A. For each set of patterns, joint detections were counted for acquisition times t_{aq} for a total measurement time

$t = Mt_{aq}$ to make up the measurement vector y. The joint distribution $p(u,v)$ was reconstructed using a Gradient Projection solver for Eq. 5 with ℓ_1 regularization, commonly referred to as Basis Pursuit Denoising [87].

We measured at dimensions of $N = 256^2$, $N = 576^2$, and $N = 1024^2$ corresponding to DMD resolutions of 16×16, 24×24, and 32×32 pixels. The associated measurement numbers M were 2500, 10,000, and 30,000 so that M is only about $0.03N$. Acquisition times were 1 second for position measurements and 1.5 seconds for momentum measurements to average 1000 coincident detections per DMD configuration in all cases. Additionally, we performed representative simulations at 16×16 and 24×24 resolutions.

E. Joint Probability Distribution

A simulation for measuring position-position correlations at 16×16 DMD resolution is given in Fig. 8. The object in 8(a) is the correlation function of Eq. 7. The simulation used $m = 2500$ measurements and a photon flux of $\Phi = 5000$ photons/measurement multiplied by the ideal $p(u,v)$, conditions representative of the 1 second experimental acquisitions. Note that this is the total signal strength before interacting with the sensing matrix; the mean value of the measurement vector is $\Phi/4 = 1250$ detected photons. Values of the measurement vector were Poissonian distributed to simulate the effect of shot noise.

Fig. 10(b) gives the reconstructed correlation function $p(u,v)$ between signal and idler DMD pixels. The sharply defined diagonal line shows the expected positive correlations between the two DMDs. DMD pixels are listed in column-major order. The mean squared error (MSE) for the reconstruction was 5×10^{-8}. The two-dimensional signal marginal distribution is inset, which provides an image of the signal beam. Figs. 8(c) and 8(d) sum the result along the anti-diagonal to show the correlation width σ_c. Qualitatively, the reconstruction closely resembles the original object, faltering only near the edges of the distribution where the signal falls beneath a noise floor. The reconstruction recovers $\sigma_c < 1$ pixel with negligible error.

To demonstrate the reconstruction accuracy, simulations were performed for increasing photon flux Φ
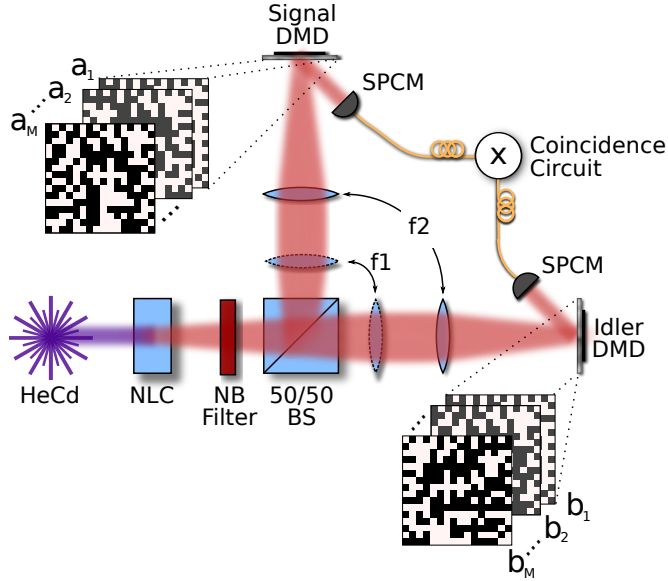

FIG. 7. Experimental Setup. Photons generated via SPDC pass through a narrow-band filter and are split into signal and idler modes by a 50/50 beamsplitter. For position correlations, lenses $f1 = 125$ mm and $f2 = 500$ mm form a $4f$ imaging system with the crystal and DMDs placed in the object and image planes respectively. For momentum correlations, f1 is removed and the DMD is placed in the focal plane of $f2 = 88.3$ mm. Photons striking DMD "on" pixels are directed to large area, single-photon counting modules. Photon arrivals are then correlated by a coincidence circuit.
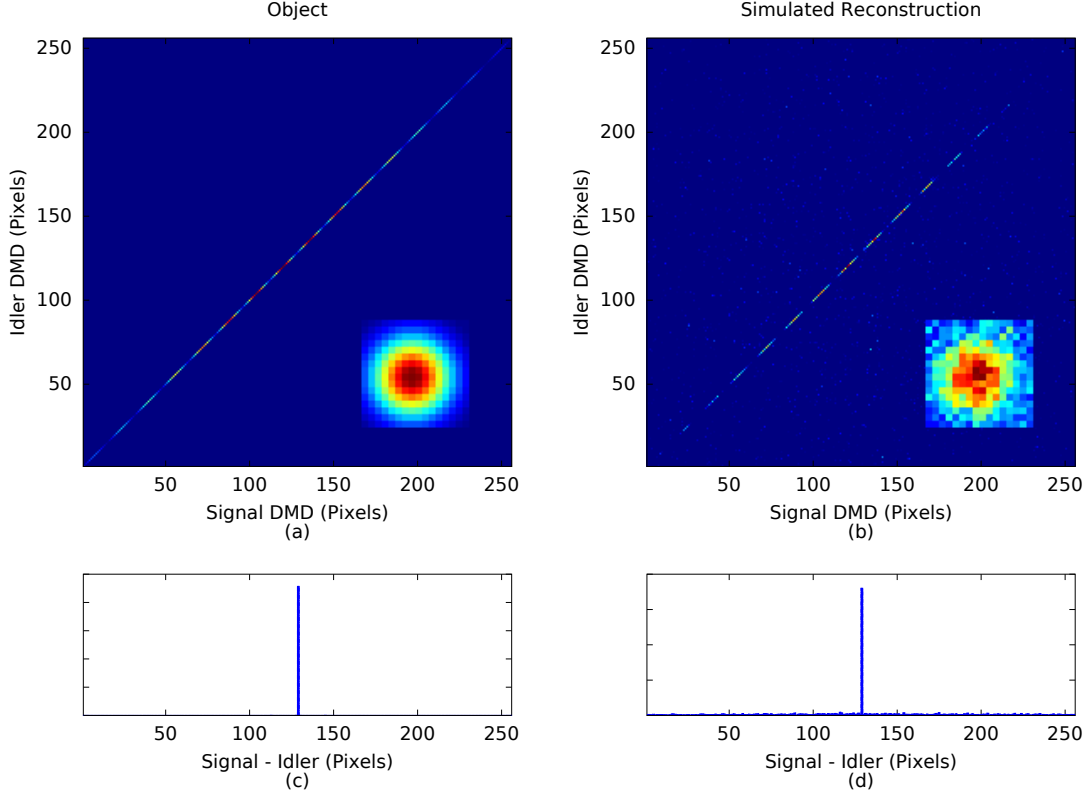
FIG. 8. 16×16 Pixel Simulation. The ideal object is given in (a). The object was incoherently sampled with $m = 2500$ random binary patterns. Poissonian noise corresponding to 5000 photons in the field (≈ 1250 detected) per measurement was added to the measurement vector. The reconstruction is shown in (b), with MSE 5×10^{-8}. The inset images in (a) and (b) show the signal photon's 2D marginal distribution, giving an image of the signal photon. Plots (c) and (d) integrate along the anti-diagonal to show that the reconstruction recovers the correlation width $\sigma_c < 1$ pixel with negligible error.

with DMD resolution 16×16 and $m = 2500$. The MSE versus Φ is given in Fig. 9. Reconstructions were normalized to the incident flux Φ for comparison to the ideal signal. The result shows the rapid phase change from poor to excellent reconstructions with a MSE converging to 5×10^{-8} beyond the phase change.

The MSE can be used to roughly estimate the signal-to-noise ratio for a particular measurement of an average, non-zero element. Assuming perfect pixel correlations and uniform marginal distributions, the energy in the signal is distributed over $1/n$ elements. The signal-to-noise ratio is then $1/n\sqrt{\text{MSE}}$. For $n = 256$ pixels and MSE $= 5 \times 10^{-8}$, this yields an approximate SNR of 17. For comparison, using Eq. 12, a raster scan would require about four days to achieve a SNR of only 10. The simulated CS acquisition time was 42 minutes for 2500, 1 second measurements.

Sample experimental reconstructions for position-position and momentum-momentum correlations at 16×16 pixel DMD resolution are given in Fig.10. As in the simulations, the position-position result shows a well defined diagonal line indicating positive pixel correlations. Conversely, the momentum-momentum result shows an anti-diagonal line showing the expected anti-correlations. Figs. 10(c) and 10(d) sum the results along the anti-diagonal (position-position) and diagonal (momentum-momentum) to reveal an effective correlation width σ_{ce} of a single pixel. Our detection scheme is therefore as accurate as possible at this resolution and our channel capacity remains detector limited.
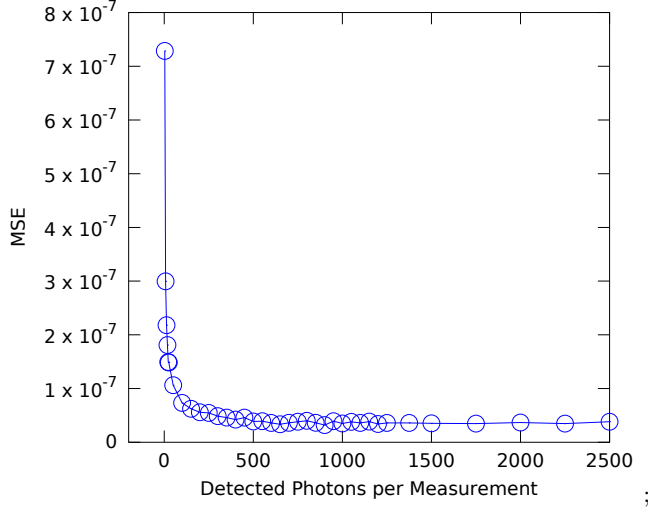
FIG. 9. Simulated Mean Squared Error (MSE) versus Photon Flux for $n = 256$ and $m = 2500$. The phase change behavior versus photon number can be clearly seen. The experiment used 5000 total (1250 detected) photons per measurement to comfortably exceed the phase change. The MSE approaches a value 5×10^{-8} corresponding to an SNR of about 17.

F. Mutual Information in the Channel

Once $p(u, v)$ is recovered, the channel capacity is given by the classical mutual information shared between signal and idler DMD pixels;

$$I = -\sum_u p(u) \log p(u) - \sum_v p(v) \log p(v)$$
$$+ \sum_{u,v} p(u, v) \log p(u, v), \tag{13}$$

where for example,

$$p(u) = \sum_v p(u, v) \tag{14}$$

is the signal particle's marginal probability distribution [28]. This entropic analysis is solely measurement based and does not require reconstructing a wavefunction or density matrix, a challenging task even for low-dimensional systems [56–58].

To estimate the uncertainty in the mutual information from shot noise and the reconstruction process, we performed 100 simulations at $n = 256$ pixel resolution and 31 simulations at $n = 576$ pixel resolution. Simulations were not performed at $n = 1024$ pixel resolution due to available computer time. In addition to the results from the raw reconstruction, thresholding was performed to provide noise reduction, where all values in the recovered $p(u, v)$ below a percentage of the maximum value are forced to zero. The simulated mutual information versus thresholding percentage is given in Fig. 11 for the $n = 256$ pixel simulations exemplified by Fig. 8. Errorbars enclose one standard deviation from repeated simulations.

As the threshold increases from zero, the mutual information rises as a weak, uncorrelated noise floor is removed. An optimal threshold is quickly reached, beyond which the threshold removes more signal than noise, reducing the mutual information. Note that the reconstructed mutual information is systematically lower than the actual mutual information in the ideal object. This is due to remaining noise and difficulty in recovering parts of the signal towards the tail of the distribution.

The $n = 256$ far field experimental result is included for comparison to the simulation. The experiment closely matches the simulation both for no thresholding and beyond its optimal threshold, but is smaller
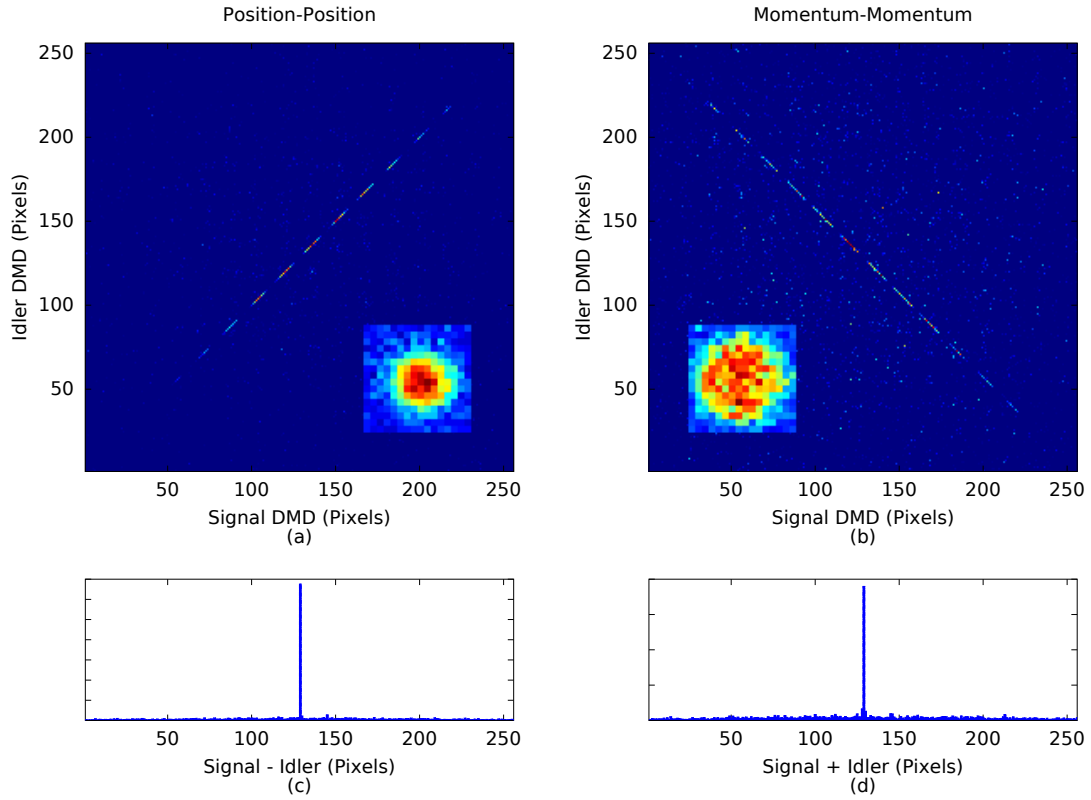
FIG. 10. Sample 16×16 Experimental Reconstructions. 10(a) and 10(b) give the joint probability distribution for position-position and momentum-momentum correlations where DMD pixels are listed in column-wise order. 2D marginal distributions for the signal photon are inset. 10(c) and 10(d) show correlation widths of only 1 pixel by summing over the signal $+$ idler (c) and signal $-$ idler (d) axes. Only 2500 (3% of raster-scanning) measurements were needed with a total acquisition time of about 40 minutes.
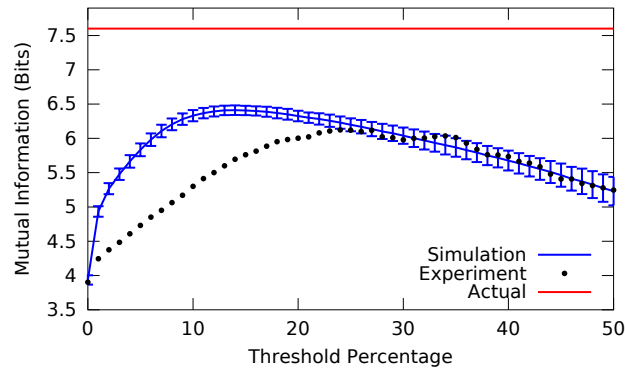

FIG. 11. Mutual Information Versus Thresholding. The mutual information for reconstruction values above a thresholded percentage of the maximum is given for 100, $n = 256$ pixel simulations with $m = 2500$ measurements and $\Phi = 5000$ photons per measurement. The red, constant line gives the true mutual information for the simulated object. The black points give the $n = 256$ far field experimental data for comparison.
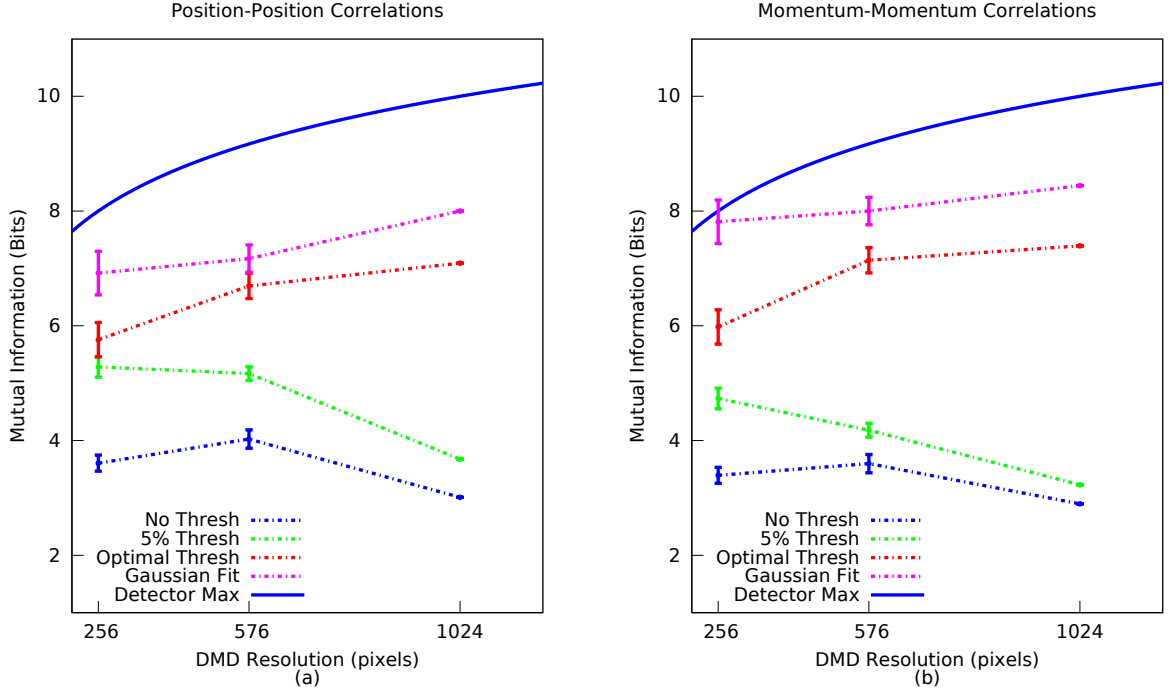
FIG. 12. Mutual information between signal and idler photons for position-position and momentum-momentum representations are presented as a function of detector resolution. Three levels of thresholding are shown, as well as a fit to Eq. 8. Dashed lines are to guide the eye. Error bars enclose two standard deviations from the expected uncertainty from simulations (not performed for $n = 1024$). The solid curve represents the maximum possible value for a particular detector resolution given perfect correlations and uniform marginals.

in the intermediate region. This is likely due to experimental errors not included in the simulation. These include slight pixel misalignment between signal and idler DMDs, optical aberrations, detector dark noise, stray light, power fluctuations in the laser, and temperature stability of the nonlinear crystal. Fig. 11 indicates that these experimental difficulties appear to increase the uncorrelated noise floor rather than significantly affect the correlated part of the reconstruction.

Although thresholding is a simple post-processing technique, it is applicable to how the entangled pixels might be used for communication. If a pair of entangled pixels has a correlated amplitude near or below the background noise, it will be difficult to use that particular mode for communication. Note that thresholding and similar noise reduction techniques can not be used if a communication protocol encodes information on single instances of the state. However, any entanglement characterization will necessarily require many instances, so background noise can often be removed to obtain a more accurate measurement. This is similar to the technique in photonic quantum information of subtracting background noise from a measured signal. In CS, it is common to perform post-processing or secondary optimization after maximizing sparsity, such as the debiasing routine in Ref. [87].

The experimental channel capacity versus DMD resolution for both position-position and momentum-momentum is given in Fig. 12 for several levels of thresholding. The optimal threshold is that which maximizes the mutual information. At 256 and 576 pixel resolutions, optimal thresholds of 20% and 30% were used for position-position and momentum-momentum distributions respectively. At 1024 pixel resolution, noise was more significant, so the optimal thresholds increased to 30% and 40%. Error bars on $n = 256$ and $n = 576$ pixels measurements represent the expected effect of shot noise and reconstruction uncertainty derived from simulation. These have been conservatively set to include two standard deviations from the simulated result.

The joint probability distribution was also fit to the double-Gaussian wavefunction (Eq. 8) to find effective
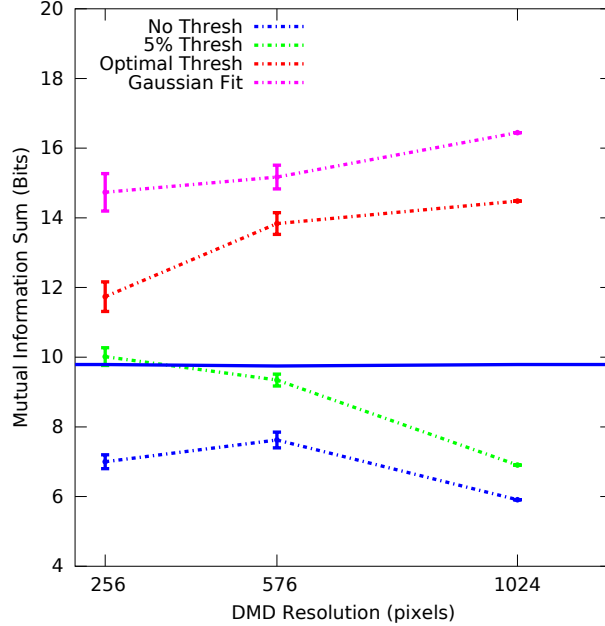
FIG. 13. The sum of position-position and momentum-momentum mutual information is presented as function of detector resolution to demonstrate violation of an EPR steering inequality (Eq. 16). The solid line represents the threshold that must be *exceeded* to witness EPR steering. Error bars enclose two standard deviations from the expected uncertainty from simulations (not performed for $n = 1024$). Because simulations systematically underestimate the mutual information, values above the bound are convincing.

widths σ_{ce} and σ_{pe}. When $\sigma_p \gg \sigma_c$, the mutual information between particles for Eq. 8 is the logarithm of the Federov ratio [59]

$$\log\left(\frac{\sigma_p^2}{\sigma_c^2}\right), \tag{15}$$

where the ratio is squared for two dimensions. While this technically applies to the continuous wavefunction, and the true σ_c is smaller than a DMD pixel, Eq. 15 still applies to the discretized measurement so long as the effective $\sigma_{ce} \gg \sigma_{pe}$.

Fitting yielded the largest channel capacities with a maximum of 8.4 bits for momentum-momentum correlations at 1024 pixel resolution, equivalent to 337 independent, identically distributed, entangled modes.

Given that fitting more accurately characterizes the system and gives a larger mutual information, it is reasonable to question the usefulness the direct mutual information computation. However, the two approaches suit different purposes. Fitting is useful if one is particularly interested in the state itself. However, if one intends to use correlated pixels for some other purpose, such as communication, the direct calculation is more appropriate. This is because the correlated pixels on the low intensity, tail of the distribution will be difficult to use in practice even if their amplitude can be inferred by fitting.

The solid curve of Fig. 12 gives the maximum possible mutual information between two, n-pixel detectors. Assuming perfect diagonal or anti-diagonal correlations and uniform marginals, this is simply $\log(n)$. Because we have Gaussian marginals, we do not expect to reach this bound, even with $\sigma_{ce} \approx 1$ pixel. By magnifying and using only the central part of the field, we could approach this upper limit.

G. Witnessing Entanglement

Despite not reconstructing a full density matrix, it is still possible to demonstrate non-classical behavior by comparing position-position and momentum-momentum correlation measurements directly. This has traditionally involved fitting the measurements to Eq. 8 and analyzing products or sums of conditional variances [60–62].

We recently presented a more inclusive, entropic steering inequality for witnessing continuous variable entanglement with discrete measurements [30], where the sum of the classical mutual information between position-position and momentum-momentum correlations is classically bounded. For our system, all classically correlated measurements must satisfy

$$I_{xs,xi} + I_{ks,ki} \leq 2 \log\left(\frac{n d_x d_k}{\pi e}\right), \tag{16}$$

where d_x and d_k are the respective widths of DMD pixels in the position and momentum basis. Note that $n d_x d_k$ is simply the bandwidth product for the DMD area, and is independent of n if total area does not change.

The sum of the classical mutual information in conjugate bases for each detector resolution is given in Fig. 13. The solid blue line provides right hand side of Eq. 16, which must be exceeded to witness EPR steering. Error bars for the $n = 256$ and $n = 576$ cases are derived from simulation and include two standard deviations. In all cases, we show EPR steering both with optimal thresholding and fitting to the double Gaussian wavefunction (Eq. 8). Even at 5% thresholding, there is a violation for 16×16 dimensions. Recall that simulations (Fig. 9) systematically under-represented the object mutual information relative to measurement uncertainty, so measurement error is highly unlikely to have over-estimated this sum. For the fitted 32×32 dimensional result, we violate the classical bound by 6.6 bits.

III. LINEAR COMPRESSIVE DEPTH MAPPING AND OBJECT TRACKING

High-resolution three-dimensional imaging, particularly using time-of-flight (TOF) lidar, is very desirable at ultra-low light levels. Such weakly-illuminated systems are safer, require less power, and are more difficult to intercept than their high-power equivalents. Potential applications include large- and small-scale surface mapping, target recognition, object tracking, machine vision, eye-safe ranging, and atmospheric sensing [65–69].

In a TOF lidar system, a scene is illuminated by a sequence of laser pulses. The pulses reflect off targets in the scene and return to a detector. Correlating the returning signal with the outgoing pulse train provides the TOF, which is converted to distance-of-flight (DOF). Transverse spatial resolution can be obtained either by scanning through the scene pixel-by-pixel or by using a detector array, such as a time-resolving CCD.

To achieve low-light capability, many recent efforts use photon-counting detectors for the detection element(s). Such detectors provide shot-noise limited detection with single-photon sensitivity and sub-ns timing. Examples include established technologies such as photo-multiplier tubes and geiger-mode avalanche photo-diodes (APDs), as well as more experimental devices such as superconducting nano-wires [70].

While they provide exceptional sensitivity and noise performance, photon-counting systems are challenging to scale to high transverse resolution. Systems that scan a single-element detector benefit from mature technology, but suffer from acquisition times linearly proportional to the spatial resolution. Significant per-pixel dwell times limit real-time performance to low resolution.

Alternately, researchers have constructed photon counting detector arrays, such as that used by MIT Lincoln Labs JIGSAW [71, 72] and its successors [73]. Fabrication difficulties limit sensor sizes to a current-best resolution of 32×128 pixels, with 32×32 pixels available commercially [74]. These sensors suffer from high dark count rates, pixel cross-talk, and significant readout noise. For TOF ranging, each pixel must be simultaneously correlated with the pulse train, a resource intensive process.

The goal is to overcome these challenges without increasing system cost and complexity. A promising option uses a standard, single-element, photon counting detector as the detection element in a compressive sensing, single-pixel camera [75]. The single-pixel camera leverages a scene's compressibility to recover an

image from an under-sampled series of incoherent projections. This is accomplished by imaging the scene onto a sequence of psuedorandom patterns placed on a spatial light modulator (SLM) or digital micro-mirror device (DMD). The inner-product of the scene and pattern is measured by a single-element, "bucket" detector. Optimization routines then recover the image from many fewer projections than pixels, often less than 10 percent.

The single-pixel camera can be adapted for TOF depth mapping by switching to active, pulsed illumination. Standard CS techniques require measurements to be linear projections of the signal of interest. Unfortunately, in a TOF single-pixel camera, the depth map is non-linearly related to the acquired measurements, as they contain both depth and intensity information [76, 77]. Current approaches mitigate or altogether avoid this issue in various ways. In 2011, we presented a proof-of-principal lidar camera where CS is used only to acquire transverse information [78]. Range is determined by gating on a TOF of interest. This reduces to the conventional single-pixel camera problem, but requires separate reconstructions for each depth of interest. Li et. al. present a similar, gated system (GVLICS) that recovers 3D images [79]. Kirmani et. al. [77] introduce a more sophisticated protocol (CoDAC) that incorporates parametric signal processing and a reconstruction process exploiting sparsity directly in the depth map. This system uses intermediate optimization steps to keep the problem linear.

CS can also be applied directly in the time-domain, and has been proposed for non-spatially resolving lidar systems [80]. In principle, one could combine this with transverse CS for a full 3D voxel reconstruction. This signal is extremely high dimensional and presents substantial measurement and reconstruction challenges.

In this manuscript, we present an alternative single-pixel camera adaptation for photon-counting depth mapping and intensity imaging. Rather than recover the depth map directly, we separately recover an intensity map and an intensity $\times$ depth map, where the depth map is simply their ratio. Linear projections of these signals are easily acquired using a TOF single-pixel camera with a photon counting detection element, with available light levels around 0.5 picowatts. Our approach is straightforward and uses standard CS formulations and solvers. We demonstrate imaging at resolutions up to 256×256 pixels with practical acquisition times as low as 3 seconds. We also show novelty filtering; finding the difference between two instances of a scene without recovering the full scenes. Finally, we demonstrate real-time video and object tracking at 32×32 pixel resolution and 14 frames-per-second.

A. Single-pixel Camera

The quintessential application of CS is the single-pixel camera [75]. Duarte et. al. cleverly implemented an incoherent sensing matrix by imaging a scene onto a digital micro-mirror device (DMD) of the sort found in digital movie projectors. The DMD is an array individually addressable mirrors that can be selectively oriented to reflect light either to a single-pixel "bucket" detector or out of the system. The measured intensity yields the inner-product of a pattern placed on the DMD with the scene of interest. To perform the set of measurements described in Eq. 4, pseudo-random binary patterns are placed sequentially on the DMD, where each pattern represents a row of sensing matrix A. Elements of Y are the measured intensities for corresponding patterns. An image of the scene is then recovered by an algorithm that solves Eq. 5.

CS and the single-pixel camera sparked a mini-revolution in sensing as researchers found their ideas could be easily and directly applied to common problems across a breadth of fields from magnetic resonance imaging [83] to radio astronomy [84] to entanglement characterization [85, 86].

B. Adapting Single Photon Counting for Depth Mapping

The single-pixel camera can be adapted for depth mapping at very low levels by switching to active, pulsed illumination and using a photon-counting detection element. Our setup is given in Fig. 14. A scene containing targets at different depths is flood illuminated by a pulsed laser diode. The scene is imaged onto a DMD which implements an incoherent sensing matrix. Light from DMD "on" pixels is redirected to a photo-multiplier module that detects individual photon arrivals. A time-correlated single-photon counting module (TCSPC) correlates photon arrivals with the outgoing pulses to find each photon's TOF. Full experimental
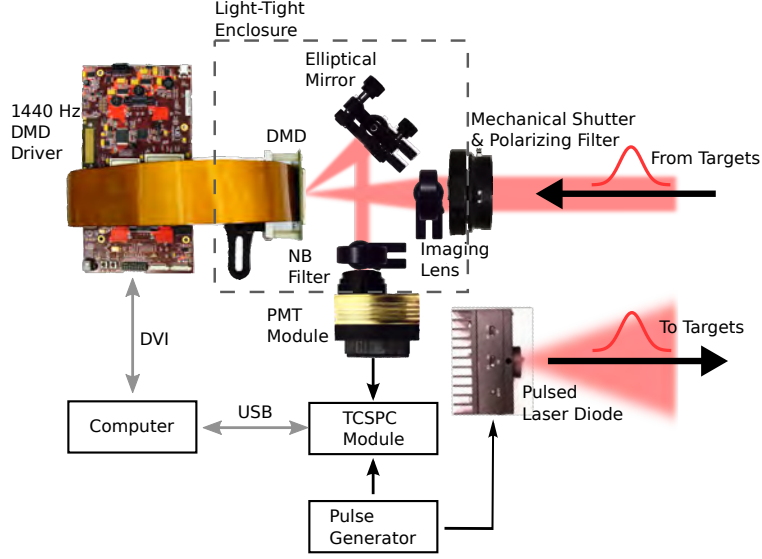
FIG. 14. **Experimental Setup** A 780 nm, 10 Mhz, 2 ns pulsed laser diode flood illuminates a scene containing targets at different depths. Returning pulses are imaged onto a DMD array with resolution up to 256×256 pixels. A polarizer prevents flares from specular reflection. Light reflecting off DMD "on" pixels is directed through a narrow-band filter to a single-photon sensitive PMT module that produces TTL pulses. Typical count rates are about 2 million photons per second. A TCSPC module time-correlates photon arrivals with the outgoing pulse train to determine a TOF for each detected photon. A series of psuedorandom, binary patterns consisting of randomly permuted, zero-shifted Hadamard patterns are placed on the DMD with per-pattern dwell times as short as 1/1440 sec. These implement an incoherent sensing matrix. For each pattern, the number of photon arrivals and the sum of their TOF's is recorded. Our protocol is then used to reconstruct the intensity image and the depth map.

specifications are given in Table I.

Recovery of a two-dimensional, intensity only image X_I (a gray-scale image) is identical to the normal single-pixel camera. Patterns from the sensing matrix are placed sequentially on the DMD. For each pattern, the number of detected photons is recorded to obtain a measurement vector Y_I. Eq. 5 is then used to find the intensity image X_I.

Unfortunately, recovering a depth map X_D from the TOF information is not straightforward because the measurements are nonlinear in X_D. Consider photon arrivals during one pattern. Each photon has a specific TOF, but is only spatially localized to within the sensing pattern. It is therefore possible and likely, that multiple photons arrive from the same location in the scene. Individual detection events therefore contain information about both intensity and depth. Consider instead a signal X_Q made up of the element-wise product $X_Q = X_I.X_D$. Unlike X_D, X_Q can be linearly sampled by summing photon arrival times for each pattern to make up a measurement vector Y_Q. Eq. 5 is therefore suitable for recovering X_Q.

The depth map is simply the element-wise division of X_Q by X_I. Note that Y_I and Y_Q are acquired simultaneously from the same signal; Y_I represents that number of photon arrivals for each pattern and Y_Q is the sum of photon TOF for each pattern. The only increase in complexity is that two optimization problems must now be solved, but these are the well understood linear problems typical for CS.

C. Protocol

A slightly more sophisticated approach including noise reduction vastly improves practical performance. The protocol we use for depth map recovery is as follows:

1. Acquire measurement vectors $Y_Q = AX_Q$ and $Y_I = AX_I$, where $X_Q = X_I.X_D$.

TABLE I. **System Specifications**

Parameter	Comments
Laser	Laser Diode mounted in Thorlabs TCLDM9
Repetition Rate	10 MHz
Pulse Width	2 ns = 60 cm
Wavelength	780 nm
Average Output Power	4 mW
Average Detected Power	0.5 pW (2 megacounts-per-second)
Spectral Filters	780/10 nm Bandpass, Optical Depth 6+
DMD Device	Digital Light Innovations D4100 with 1024×768 DMD, operated in binary expansion mode at 1440 frames per second
Spatial Resolution	images: 256×256 pixels
	video: 32×32 pixels
Detector	Horiba TBX-900C Photomultiplier Module with 9 mm^2 active area
Jitter	< 200 ps
Data Acquisition	PicoQuant PicoHarp 300 operating in time-tagging mode
Imaging Lens Focal Length	75 mm

2. Use sparsity maximation routine (Eq. 5) on Y_Q to recover noisy $\bar{X}_Q$.

3. Apply hard thresholding to $\bar{X}_Q$. The subset of non-zero coefficients in a sparse representation of $\bar{X}_Q$ now form an over-determined, dense recovery problem.

4. Perform a least squares debiasing routine on non-zero coefficients of $\bar{X}_Q$ in the sparse representation to find their correct values and recover X_Q.

5. Take significant coefficients of X_I to be identical to X_Q. Perform the same least squares debiasing on these coefficients of X_I.

6. Recover X_D by taking $\mathrm{Nz}(X_I).X_Q./X_I$, where $Nz(x) = 1$ for non-zero x and 0 otherwise.

7. (optional) In the case of very noisy measurements, perform masking on X_D and X_I.

First, we acquire measurement vectors Y_I and Y_Q as previously described. Rather than independently solve Eq. 5 for both measurement vectors, we only perform sparsity maximization on Y_Q. In practice, solvers for Eq. 5 tend to be more effective at determining significant coefficients than finding their correct values, particularly given noisy measurements [87]. We therefore rely on sparsity maximization only to determine which coefficients are significant. We subject the noisy, recovered $\bar{X}_Q$ to uniform, hard thresholding [88] in a sparse representation, which forces the majority of the coefficients to zero. The sparse basis is typically wavelets, but may simply be the pixel basis for simple images.

The subset of coefficients remaining after thresholding is considered significant. If only these coefficients are considered, the problem is now over-determined and a least-squares debiasing routine is applied to find their correct values and yield the signal X_Q[87]. We assume the significant coefficients for X_I are the same as for X_Q, and apply least-squares debiasing to X_I as well. Finally, we recover the depth map X_D by taking

$$X_D = \mathrm{Nz}(X_I).X_Q./|X_I| \tag{17}$$

where Nz$(p) = 1$ for nonzero p and 0 otherwise. This prevents a divide-by-zero situation when an element of $X_I = 0$, meaning no light came from that location.

For very noisy measurements, the sparsity maximization in steps 2-3 gives accurate outlines for targets in the scene, but poorly recovers pixel values. After the least squares debiasing in steps 4 and 5, values are more accurate, but the outlines can become distorted. A best-of-both-worlds solution is to threshold the result of step 4 in the pixel basis to create a binary object mask. This mask can optionally be applied to the final X_I and X_D for spatial clean up.

D. Experimental Setup

Our experimental setup is given in Fig. 14, with specifications given in Table I. Targets are flood illuminated with 2 ns, 780 nm laser pulses with a 10 MHz repetition rate. Average illumination power is 4 mW. The scene is imaged onto a DMD array, where on pixels are redirected through a 780/10 nm bandpass filter to a photon-counting photo-muliplier module. A sequence of m sensing patterns are displayed on the DMD at 1440 Hz. For each pattern, the number of photon arrivals and the sum of photon TOF is recorded. If 1/1440 seconds is an insufficient per-pattern dwell time t_p, the pattern sequence is repeated r times so that $t_p = r/1440$ and the total exposure time t is $t = mt_p = mr/1440$. The average detected photon rate is 2 million counts-per-second, or 0.5 picowatts.

Sensing vectors (patterns) are randomly selected rows of a zero-shifted, randomly-permuted Hadamard matrix. The Hadamard transform matrix is closely related to the Fourier transform matrix, but entries only take real values 1 or -1. The zero shifted version sets all -1 values to zero. When randomly permuted, they make practical sensing vectors because repeated calculations of AX in the reconstruction algorithms can be computed via fast transform. This is computationally efficient because the fast transform scales logarithmically and the sensing matrix does not need to be stored in memory [89].

The sparsity promoting regularization is either the signal's ℓ_1 norm or total variation depending on the scene and exposure time. For ℓ_1 minimization, we use a gradient projection solver [87]. For TV minimization,
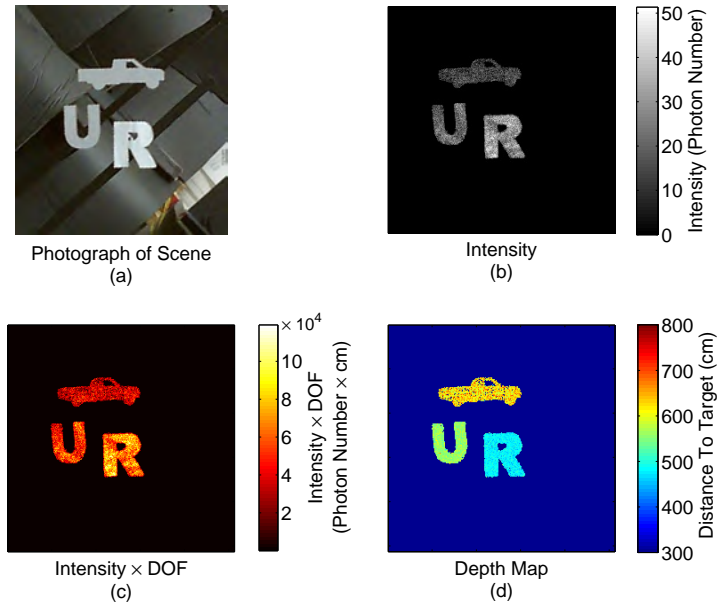

FIG. 15. **Long Exposure** A scene consisting of cardboard cutouts (a) is imaged at $n = 256 \times 256$ resolution from $m = 0.2n$ random projections with a 6.07 minute exposure time. The protocol recovers a high quality intensity map (b) and intensity $\times$ DOF map (c). Their ratio is the depth map (d).

23 Second Exposure — Intensity Map (a)
23 Second Exposure — Depth Map (b)
3 Second Exposure — Intensity Map (c)
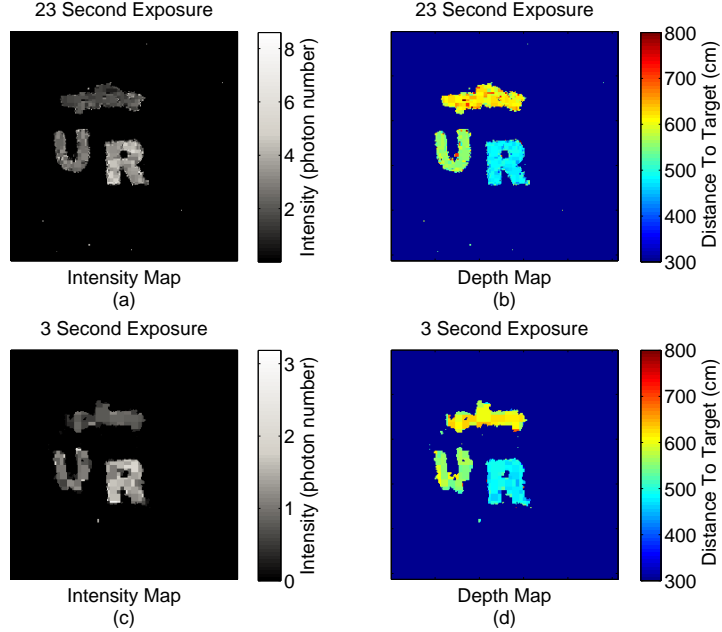3 Second Exposure — Depth Map (d)

FIG. 16. **Short Exposure** Scene from Fig. 15 acquired with rapid acquisition times. The intensity and depth maps in (a) and (b) were acquired in 23 seconds while (c) and (d) were acquired in only 3 seconds. At these exposure times, the intensity map regularly contains less than one photon per significant pixel, so it is impossible to raster scan.

We use the TVAL3 solver [89]. The sparse representation for wavelet thresholding and least-squares debiasing is Haar Wavelets.

E.   Simple Scene

The first test case is a simple scene containing cardboard cutouts of the letters "U" and "R" as well as the outline of a truck. The objects were placed at to-target distances of 480 cm , 550 cm, and 620 cm respectively. Fig. 15 shows a high quality reconstruction with resolution $n = 256 \times 256$ pixels, $m = 0.2n = 13108$, and $t_p = 40/1440$ for an exposure time of 6.07 minutes. The sparsity promoter was TV. Both the intensity and depth map are accurately recovered.

Pictures of the same scene with much shorter acquisition times are given in Fig. 16. Fig. 16(a) and 16(b) show the intensity and depth map for $m = 0.1n = 6554$ and $t_p = 5/1440$ seconds for a total exposure time of 23 seconds. Fig. 16(a) and 16(b) gives results for a exposure with $m = .07n = 4558$ and $t_p = 1/1440$ seconds for a total exposure time of 3 seconds. The optional masking process (protocol step 7) was used. While the reconstructions are considerably noisier than the long exposure case, the objects are recognizable and the average depths are accurate.

Note that the shortest dwell-time per pattern for the DMD is 1/1440 sec. To raster scan at this speed for $n = 256 \times 256$ pixels would require 46 seconds, so both results in Fig. 16 are recovered faster than it is possible to raster scan. If the same dwell times were used ($t_p = 5/1440$ and $t_p = 1/1440$ seconds respectively), the raster scan would take about 228 seconds and 46 seconds. Note that many significant pixels in the recovered intensity images in Fig. 16 have less than 1 photon-per-pixel. Therefore, even the longer raster scans would struggle to produce a good picture because they cannot resolve fewer than 1 photon per pixel. Our protocol is more efficient because each incoherent projection has flux contributed from many pixels at once, measuring half the available light on average.
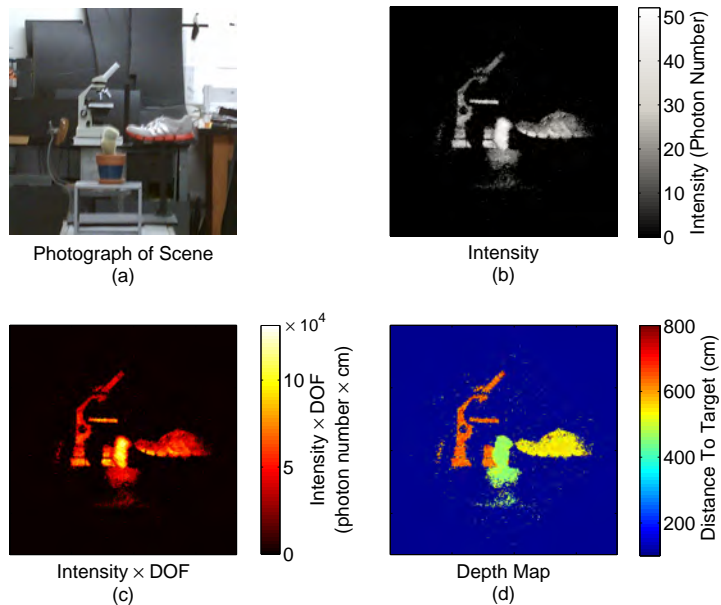
FIG. 17. **Natural Scene** A scene consisting of a cactus, shoe, and microscope is imaged at $n = 256 \times 256$ pixel resolution with $m = 0.3n$ and a per-pattern dwell time $t_p = 50/1440$ seconds.

F.   Natural Scene

A $n = 256 \times 256$ pixel measurement of a more complicated scene containing real objects is given in Fig. 17. The scene consists of a small cactus, a shoe, and a microscope placed at to-target distances of 465 cm, 540 cm, and 640 cm respectively. A photograph of the scene is given in Fig. 17(a). The scene was acquired with $m = 0.3n$ a per-pattern dwell time $t_p = 50/1440$ seconds for an exposure time of 11.37 minutes. The sparsity promoter was the ℓ_1 norm in the Haar wavelet representation.

Fig. 17(b), Fig. 17(c), and Fig. 17(d) show the reconstructed intensity, intensity $\times$ DOF, and depth map respectively. All objects are recognizable and the to-target distances are accurately recovered. Noise is slightly higher than the simpler scene of cardboard cutouts.

G.   Depth Resolution

The photon TOF measured by the TCSPC module is trivially converted to DOF by multiplying by speed of light c. This DOF includes the round trip distance traveled by illuminating pulses as well as delays introduced by cables and electronics. The actual distance to target is linearly related to the measured DOF.

To determine this relationship and probe the system's depth resolution, we took images of a simple square, white, cardboard cutout placed a known distance from the camera. The cutout was sequentially moved away from the camera in small increments and the reconstructed depths were recorded.

Results are given in Fig. 18. Fig. 18(a) shows a photograph of the object, with a typical reconstructed depth map given in Fig. 18(b). Fig. 18(c) shows a coarse grained result when the object was moved in 15.25 cm (6 in) increments over a 259 cm (108 in) range. Each measurement was performed at 32×32 pixel transverse resolution with $m = 0.1n = 102$. The per-pattern dwell time was $1/1440$ seconds, so each depth map was acquired in .07 seconds. The pulse length was 2 ns, or 60 cm.

Fig. 18(c) provides the reconstructed DOF as function of the distance to object. The measured DOF is averaged over the recovered object. A linear fit shows very good agreement with the measurements, with a slope of 1.91. This is slightly less than the expected 2 (for a round-trip flight) because the result includes
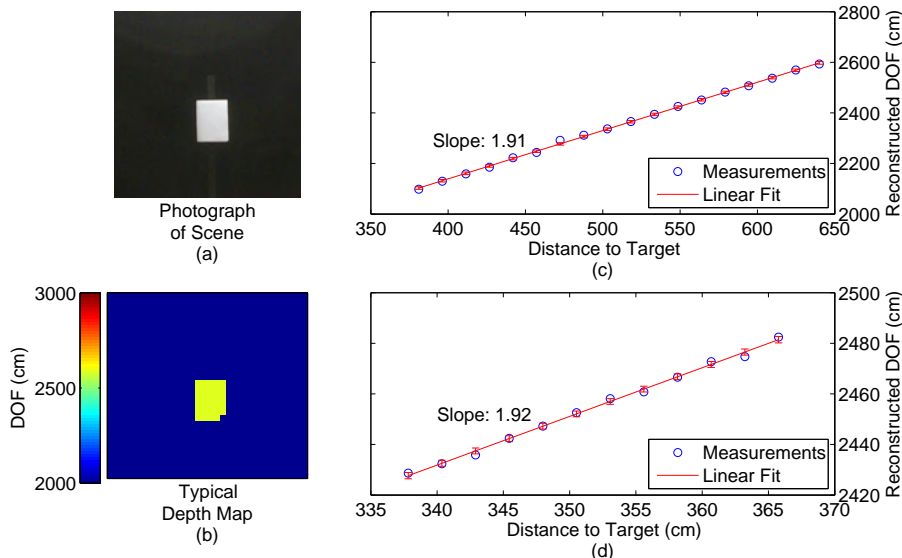
FIG. 18. **Depth Calibration** Depth maps of a rectangular cardboard cutout (a) are acquired at $n = 32 \times 32$ pixel resolution with a typical reconstruction given in (b). The cutouts to-target distance was increased in increments of 15.52 cm (c) and 2.54 cm (d). Depths can accurately recovered to less than 2.54 cm for this scene.

cable transit times. Fit error bars represent a 95% confidence interval.

We performed an additional fine resolution set of measurements where the object was moved in 2.54 cm (1 in) increments over a 30.4 cm (12 in) range. All other parameters are identical to the coarse measurements. Results are shown in Fig. 18(d). Despite 60 cm pulse lengths, depth mapping is accurate to less than 2.54 cm. While uncertainty is slightly larger than the coarse case, the fits agree to within the confidence interval.

H. Novelty Filtering

For many applications such as target localization, high frame-rate video, and surveillance, it is useful to look at the difference between realizations of a scene at different times. Such novelty filtering removes static clutter to reveal changes in the scene. This is traditionally accomplished by measuring the current signal $X^{(c)}$ and subtracting from it a previously acquired reference signal $X^{(r)}$ to find $\Delta X = X^{(c)} - X^{(r)}$.

Compressive sensing is particularly well suited for novelty filtering because the difference signal ΔX can be directly reconstructed [90, 91]. Consider acquiring incoherent measurement vectors for both $X^{(r)}$ and $X^{(c)}$ as in Eq. 4, using the same sensing matrix A for both signals. Instead of separately solving Eq. 5 for each signal, first difference their measurement vectors to find

$$\Delta Y = Y^{(c)} - Y^{(r)} \tag{18}$$

$$\Delta Y = A\Delta X. \tag{19}$$

Because Eq. 19 has the same form as Eq. 4, the optimization routine can be performed directly on ΔY to obtain ΔX without ever finding $X^{(c)}$ or $X^{(r)}$. Furthermore, the requisite number of measurements m depends only on the sparsity of ΔX. For small changes in a very cluttered scene, the change can be much sparser than the full scene. It is often possible to recover the novelty with too few measurements to reconstruct the full scene.

Fig. 19 gives an example of using our system for novelty filtering in intensity and depth. A photograph of a current scene containing several objects is shown in 19(a). Fig. 19(b) photographs a prior reference scene.
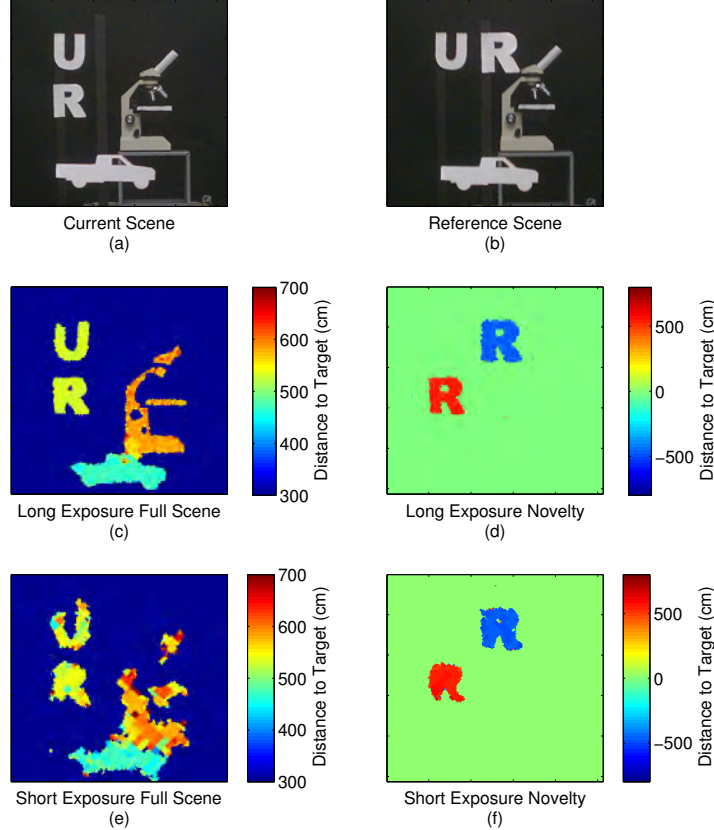
FIG. 19. **Novelty Filtering** (a) and (b) give photographs of two instances of a scene, where the 'R' has changed positions (including depth). (c) and (d) show high quality, long 6.07 minute exposure depth map reconstructions for the full current scene and difference image respectively. (e) and (f) show corresponding short 37 second exposure depth map reconstructions. Negative values in the difference image indicate the object's former location.

In the current scene, the cardboard cutout 'R' has been moved, both transversely and from a range of 480 cm to a range of 560 cm.

Fig. 19(c) and 19(d) give long exposure reconstructions of the full current depth map and the difference depth map respectively. Exposure parameters were $n = 256 \times 256$ pixels, $m = 0.2n$, and $t_p = 40/1440$ sec for an exposure time $t = 6.07$ minutes. The sparsity promoter was TV. The difference reconstruction contains only the 'R' object that moved. Note that there are two copies of the 'R'. One is negative, indicating the objects former position, and one is positive, indicating the objects new position.

Fig. 19(e) and 19(f) show short exposure reconstructions for the current scene and the difference image. The number of measurements was reduced to $m = 0.02n = 1311$ with $t_p = 40/1440$ for an exposure time of 37 seconds. Masking was performed. For this acquisition time, the static clutter is very poorly imaged and is difficult to recognize in the full scene. The difference image effectively removes the clutter and gives a recognizable reconstruction of the novelty. Again, this is far faster than raster-scanning, which requires at least two 45 second scans for differencing.

I. Real-time Video and Object Tracking

At lower resolutions, the system is capable of real-time video and object tracking. To demonstrate this, we acquired video at 32×32 pixel resolution of a three-dimensional pendulum consisting of a baseball suspended from the ceiling by a rope. The lever arm was 170 cm and the pendulum oscillated through a solid angle of

approximately 25 degrees. The three dimensions are not oscillating in phase, so the trajectory is elliptical.

Movie frames were acquired with $m = 0.1n = 99$ with a per-pattern dwell time $t_p = 1/1440$ sec. Each frame required .07 sec to acquire for a frame rate slightly exceeding 14 frames-per-second. The sparsity promoter was TV. Sample frames showing the pendulum moving through a single period are given in Fig. 20, where alternate frames are skipped. We clearly recover images of the pendulum oscillating in all three dimensions. The full movie can be viewed online.

The pendulum's position can be tracked by taking expected values for its transverse location and averaging over its range to determine a to-target distance. Fig. 21 shows the computed pendulum trajectory. Figs 21(a), 21(b) and 21(c) give the expected x, y and z values for pendulum location as a function of frame. These can be combined to yield the three-dimensional, parametric trajectory given in Fig. 21(d). Sinusoidal fits are in good agreement with the expected values, particularly in the depth dimension.

IV.   GHOST IMAGING AND POINT TARGET LOCALIZATION

Ghost images are obtained by correlating the output of a single-pixel (bucket) photodetector—which collects light that has been transmitted through or reflected from an object—with the output from a high spatial-resolution scanning photodetector or photodetector array whose illumination has not interacted with that object. The term "ghost image" is apt because neither detector's output alone can yield an image: the bucket detector has no spatial resolution, while the high spatial-resolution detector has not viewed the object. The first ghost imaging experiment relied on the entangled signal and idler outputs from a spontaneous parametric downconverter [92], and hence the image was interpreted as a quantum phenomenon. Subsequent theory and experiments showed, however, that ghost images can be formed with pseudothermal light [93, 94]. A controversy soon arose as to whether a quantum interpretation was required for understanding the pseudothermal case [95, 96]. This controversy, since resolved [97–99], obscured the more significant issue insofar as standoff sensing is concerned, viz., does any form of ghost imaging offer intrinsic advantages in comparison to conventional laser radar imaging? In [100], described below, we completed the first comparison of the spatial resolutions and signal-to-noise ratios for such systems when target speckle and atmospheric turbulence are accounted for.

Ghost imaging—especially its computational form [101], as considered in [100]—is more properly termed structured-illumination imaging, for which image-formation algorithms far more efficient than correlation have been demonstrated [102]. Nevertheless, structured illumination's ultimate photon efficiency for standoff sensing against speckle targets in the presence of atmospheric turbulence is unknown. In ongoing work [103], described below, we are developing an appropriate version of generalized approximate message-passing (GAMP) to substantially increase structured illumination's photon efficiency in this scenario.

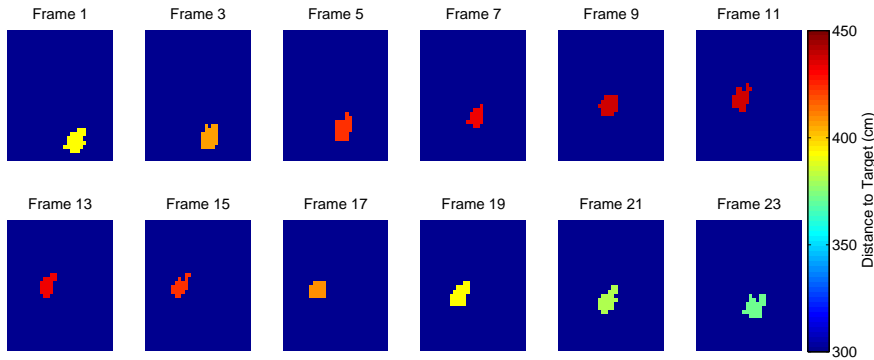Laser radars offer superior resolution in comparison with their microwave counterparts. Laser radar,

FIG. 20. **Movie** Frames from a depth-map movie of a three-dimensional pendulum consisting of a baseball suspended by a 170 cm rope swinging through a 25 degree solid angle. The transverse resolution is 32×32 pixels with a frame rate of 14 frames per second.
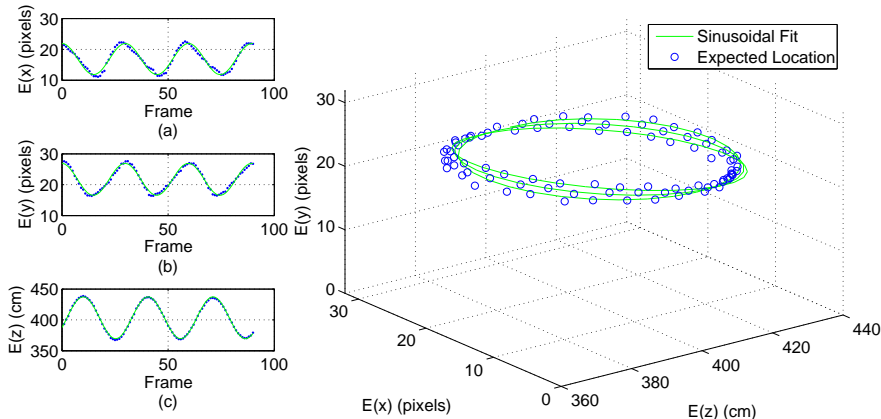
FIG. 21. **Object Tracking** The expected values for transverse (x,y) and range (z) coordinates are given in (a), (b), and (c) as a function of frame number. Blue circles show expected values obtained from reconstructed depth maps, while green lines give sinusoidal fits. (d) shows a 3D parametric plot of the pendulum's trajectory.

however, suffers from the ill-effects of target-induced speckle and atmospheric propagation. Moreover, laser radars superior resolution capability can dictate a longer time to interrogate a particular target region. Thus a useful scenario to consider is for a high-resolution laser radar to be cued by a lower-resolution microwave system. An interesting question that arises is determining the ultimate sensitivity when quantum-mechanical resources—nonclassical transmitter states and nonstandard reception techniques—are provided. In [104] we addressed that problem. As detailed below, we showed that nonclassical resources are *not* of value in this application and that high photon efficiency (multiple bits/detected-photon) is possible for speckle targets even in the presence of background light. Initial experiments for dark-count limited operation—performed by Prof. Howell's group at the University of Rochester and reported in [104]—were in excellent agreement with our theoretical predictions.

The original context for our InPho program was quantum imaging. Yet, as our team's work on point-target localization and first-photon imaging has shown, nonclassical resources have yet to prove themselves to be of particular value for high photon-efficiency standoff imaging. Thus, in ancillary work we clarified some issues with respect to the distinction between quantum and quantum-mimetic ghost imaging. First, we performed experiments comparing classical ghost images obtained, for which there is a phase-insensitive cross correlation between the signal and reference beams, with those obtained when there is a phase-sensitive cross correlation between those beams [105]. As explained below, this experiment verified theoretical predictions from [96], in which classical phase-sensitive ghost imaging was shown to mimic quantum ghost imaging in all features other than DC-coupled image contrast. Second, we addressed the role of quantum discord [106] in pseudothermal ghost imaging. Ragy and Adesso [107] had shown that quantum discord seemingly played a role in such an imager, despite its statistics being fully explained by semiclassical photodetection. We demonstrated [108] that there was no such role for quantum discord in spatial light modulator (SLM) ghost imaging [101], casting doubt on whether there was any necessity for quantum discord in ghost-image formation.

The sections that follow provide additional details for some of the InPho program accomplishments that we have summarized above.

A. Ghost Imaging versus Laser Radar for 3D Imaging

Prior to our work, there had only been cursory examinations of the relative merits of ghost imaging and conventional laser radars for standoff sensing of real-world objects. In [100] we remedied that situation. There, we presented results for the spatial resolutions and signal-to-noise ratios (SNRs) when these sensors

were used to image rough-surfaced targets that produce fully developed laser speckle and the propagation to and from the targets was through atmospheric turbulence which was distributed arbitrarily along the propagation paths. We also investigated the tradeoff between spatial resolution and SNR as a function of detector and entrance-pupil sizes. Figure 22 shows the configurations that we considered: a computational ghost imager and a direct-detection laser radar. The principal conclusions we obtained from analyzing these two configurations were as follows.
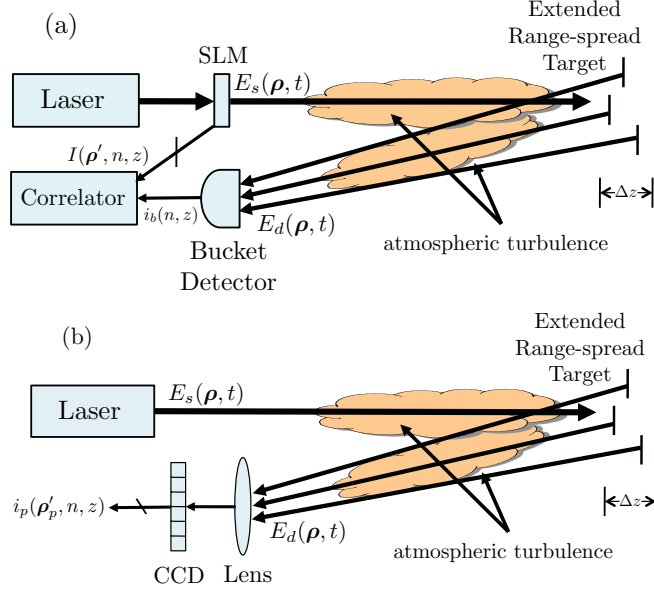

FIG. 22. (a) Setup for 3D computational ghost imaging in reflection. Pulsed laser-light undergoes spatial light modulation, propagation through atmospheric turbulence to and from an extended range-spread target, and shot-noise limited bucket detection, producing an output, $i_b(n,z)$, from the range-z return associated with the nth transmitted pulse. Diffraction theory is used to calculate $I(',n,z)$, the vacuum-propagation, target-region intensity pattern at transverse coordinate $'$ and range z associated with the nth transmitted spatial pattern. (b) Setup for 3D-imaging laser radar. Pulsed laser-light illuminates the target region, in a floodlight manner, through a turbulent optical path. The return light, collected after propagation back to the radar through atmospheric turbulence, is detected by a shot-noise limited CCD array in an image plane, resulting in an output, $i_p(',n,z)$, for the pixel at transverse coordinate $'_p$ and range z from the nth transmitted pulse.

Computational ghost imaging is, in many respects, a dual of floodlight-illumination laser radar. The computational ghost imager's system complexity is in its transmitter, whose sequence of SLM patterns creates the structured illumination that provides the imager's spatial resolution. The laser radars complexity lies in its CCD receiver, which provides its spatial resolution. Consequently, the size of the ghost imager's transmitter pupil sets its diffraction-limited spatial resolution, whereas the laser radar's no-turbulence spatial resolution is set by the size of its receiver pupil in conjunction with that of its CCD pixels. Thus, only turbulence on the transmitter-to-target path can impair the ghost imager's spatial resolution, while the laser radar's spatial resolution is only impacted by turbulence on the target-to-receiver path. Therefore, ghost imaging and laser radar systems that are designed to have equal spatial resolutions in the absence of turbulence could have significantly different performance in a bistatic configuration, in which their transmitters and receivers are not colocated, so that significantly different turbulence distributions are encountered on these two paths. In such situations, either one could offer the better spatial-resolution performance, depending on which path had its turbulence concentrated near the resolution-controlling pupil. Aside from this turbulence issue, our analysis indicated that a fair comparison between the computational ghost imager and the floodlight-illumination laser radar shows them to have equal spatial resolutions, except for the following two caveats: (1) the laser radar can form its image from a single pulse, making it far better for imaging moving targets; and (2) the computational ghost imager has infinite depth of focus, whereas the laser radar will not

for ranges in the near field of its receiver's entrance pupil.

Both the ghost imager and the laser radar have SNRs that are shot-noise limited at low photon fluxes and that saturate at high photon fluxes. When their optics are sized for equal spatial resolutions, there is little difference in their saturation SNRs. This contrasts strongly with their shot-noise-limited SNR behavior, in which the laser radar outperforms the ghost imager by a factor approximately equal to the number of spatial-resolution cells on the target. As a result, we can expect that the ghost imager will require significantly more time to achieve a desired SNR when operating in this regime. This key disadvantage for correlation-based ghost imaging could be mitigated to some degree, however, by the use of compressed-sensing techniques, which enable many fewer pulses to suffice for ghost-image formation. Recent work has demonstrated this possibility in a standoff scenario [109].

What then are the possible advantages of ghost imaging in comparison with laser radar? The principal such advantage identified by our analysis accrues in bistatic configurations wherein, for operational reasons, the transmitter must be located in a region of weak turbulence, but the receiver necessarily is in a strongly turbulent region. Beyond that, however, there are some technological possibilities. The ghost imager only requires a single-pixel detector, whereas the laser radar needs a detector array. For wavelength regions in which high-performance single-pixel detectors are available but similar-quality detector arrays are not, ghost imagers would provide active-imaging capability that laser radars could not. A related technological advantage arises for ghost imaging in multistatic configurations in which a network of simple, small, single-pixel detectors view a target region that is floodlit by a single, large-aperture, structured-illumination transmitter. Individual images could be formed from each detectors outputs to capture multiple views of the target, and allow for more averaging of the target-induced speckle. A corresponding multistatic laser radar would require high-resolution CCDs at each receiver location, making it more complicated and more expensive than the ghost imager.

B. Graphical Modeling for Active Imaging of Speckle Targets

Graphical modeling is the mapping of probabilistic relations to a structure of nodes and edges, generally as a means of defining structure that can be used for inference [110]. This structure is related to conditional dependencies, and provides a means for using local calculations to perform inference on the whole graph. There are three main types of graphical structures—directed [111], undirected [112], and factor graphs [113]—and each have scenarios in which they are preferable, but for our needs factor graphs seem most appropriate, motivated by desire to use approximate message passing (AMP) techniques [114].

We have been developing [103] a graphical-model based loopy belief-propagation (LBP) approach for optimizing a standoff active-imager that is allowed to employ structured illumination as well as a spatially-resolving detector. The field at the detector plane is a projection of the interrogating signal field onto the complex-field reflectance of the target; messages that approximate LBP can be found for these interactions using generalized approximate message passing (GAMP) [115]. Since the target is assumed to be rough, with microscopic height variations whose correlation length is on the order of a wavelength, the measurements will be speckled [117], and not an exact projection, once the target and field are discretized. However, these nonidealities can be handled with the inclusion of uncertainty in the measurement matrix (illumination pattern), as this is essentially a quantization error, and has been addressed in [116].

This work is still in progress. Very preliminary results for a single-detector system, shown in Fig. 23, indicate that employing the matrix-uncertainty generalized approximate message passing (MU-GAMP) algorithm does not produce high-quality imagery. Note that the reconstruction shown in that figure did not use a sparsifying prior, whose inclusion would certainly reduce its speckled appearance. Likewise, multiple-detector operation will further improve the results. An initial sense of the improvements that might be obtained is given by Fig. 24. The panel on the left shows the MU-GAMP reconstruction based on 10 detectors, while the panel on the right shows the reconstruction from our new hybrid-GAMP algorithm [103] using the same 10-detector data. Although the imagery in Fig. 24 is still far from high quality, it is important to note that a correlation-based imager using the same data would yield an image resembling white noise, because the number of realizations used to obtain the Fig. 24 results is far fewer than what is necessary for correlation processing to yield even an image outline.
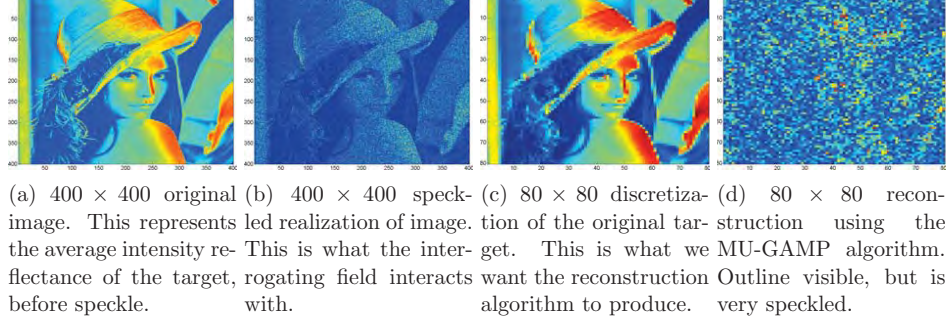
(a) 400×400 original image. This represents the average intensity reflectance of the target, before speckle.

(b) 400×400 speckled realization of image. This is what the interrogating field interacts with.

(c) 80×80 discretization of the original target. This is what we want the reconstruction algorithm to produce.

(d) 80×80 reconstruction using the MU-GAMP algorithm. Outline visible, but is very speckled.

FIG. 23. Reconstruction of an example target using patterned illumination and the MU-GAMP algorithm.
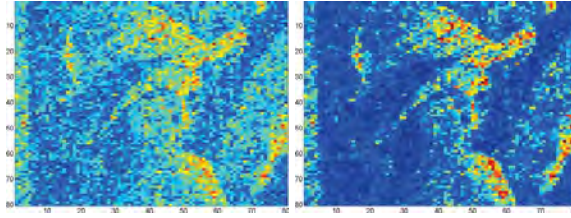
FIG. 24. Reconstruction of an example target using patterned illumination and 10 detectors with (a) the MU-GAMP algorithm, and (b) the hybrid-GAMP algorithm.

C. Point-Target Localization with High Photon Efficiency

In our quest for task-specific imaging at high photon efficiency, we considered the following point-target localization problem [104]. A pulsed-laser transmitter floodlights a volume known to contain a point target, and an M_s-pixel photon-number resolving array detects the returned light. The localization task is to determine the target's transverse location within these M_s pixels and its range within M_r range bins. In the absence of background light and detector dark counts, this scenario can yield an extraordinary number of bits per detected photon (bpdp), viz., a 32×32 pixel array combined with 15 cm range resolution and a 1 km uncertainty in target range will yield $\log_2(M_s M_r) = 22.7$ bits from the detection of one photon. There is still, however, the possibility of no detections, even though the target is present. Forcing the radar to randomly choose among the $M = M_s M_r$ possible target locations when no photons are detected then reduces bpdp to 3.3 when the error probability—$\Pr(e) = (M - 1)\exp(-\eta n_T)/M$, where n_T is the average number of transmitted photons and η is the radar-to-target-to-radar transmissivity—is 10^{-3}. This performance is realized with $n_S \equiv \eta n_T = 6.91$. With a number-state transmitter and the optimum quantum receiver, we have [118]

$$\Pr(e) = \frac{M-1}{M^2}\{[1 + (M-1)p]^{1/2} - (1-p)^{1/2}\}^2, \text{ with } p = (1-\eta)^{n_T}, \tag{20}$$

yielding $\Pr(e) = 10^{-3}$ at $\eta = 10^{-4}$ and $n_T = 6.88/\eta$, showing that little is to be gained from the nonclassical transmitter and nonstandard receiver in this case. Thus we focused on the effects of dark counts, background light, and speckle on point-target localization with a floodlight laser-transmitter and a detector array.

Our calculations were for table-top experiments aimed at verifying the theory. We assumed a 50 ps transmitter pulse duration, a 32×32 detector array with 100 dark-counts/sec on each array element, and a 50 cm uncertainty in target range. Figure 25(a) shows the erasure (no detections) and error (incorrect localization) probabilities versus n_S when dark counts are the only nonideality, and Fig. 25(c) shows these probabilities when there is also 9.4×10^3 background-counts/sec on each array element, arising from a

$1\,\mathrm{W/m^2}$-Sr-μm background spectral radiance, 1 nm optical bandwidth, and 50% detector quantum efficiency. Background counts drive up the error probability, but have little effect on the erasure probability because that is dominated by the probability that no target-return photons are detected. Figures 25(b) and (d) show the mutual information and bpdp for the dark-counts only and dark-counts plus background-counts cases, respectively. Both situations can provide more than 2 bpdp with $\Pr(\text{erasure}) \leq 10^{-3}$ and $\Pr(\text{error}) \leq 10^{-3}$.

Figure 26 shows Pr(erase) and Pr(error) [in (a)] and the mutual information and bpdp [in (b)] when, in addition to dark and background counts, the target produces fully-developed speckle. Here we see a substantial performance degradation has been incurred, but 2 bpdp can still be obtained.
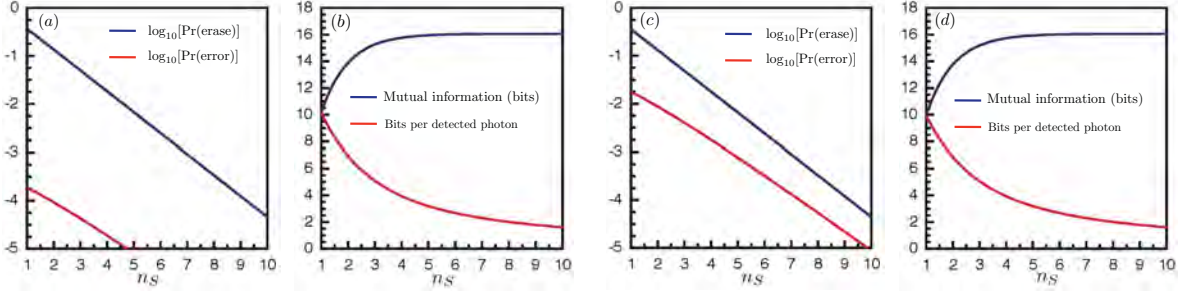

FIG. 25. Pr(erase) and Pr(error) [(a) and (c)] and mutual information (MI) and bpdp [(b) and (d)] for dark-count limited operation [(a) and (b)] and background-limited operation [(c) and (d)].
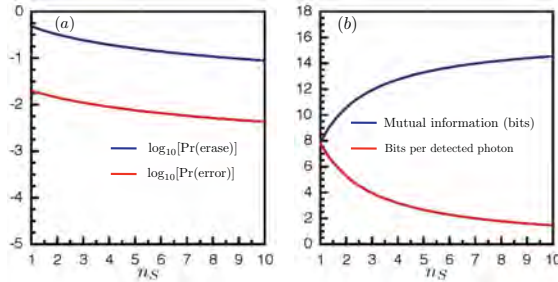

FIG. 26. Pr(erase) and Pr(error) [in (a)] and MI and bpdp [in (b)] for background-limited operation with speckle.

In a preliminary experiment, we have used the Fig. 27 arrangement to emulate point-target localization. Light from a low power HeNe laser is attenuated to the single-photon level with neutral-density (ND) filters and coupled into a single-mode fiber for spatial filtering. The fiber's output illuminates a one-to-one imaging system with digital micro-mirror devices (DMDs) in the object and image planes. The first DMD introduces a point target into the system, and the second emulates a detector array by scanning the light from its elements onto a geiger-mode avalanche photodiode.

Figure 28 compares results from an experiment in which the target is to be localized within a 16×16 pixel ($M = 256$) array in dark-count limited operation (600 dark-counts/sec on each pixel): (a) and (b) show the performance when a single pulse interrogates each pixel; (c) and (d) show the performance when each pixel is interrogated until at least one count occurs. This pulse-until-detect (PUD) protocol completely suppresses erasures, but at the expense of more errors. Figure 28 shows that this error-probability penalty is not severe, and that our experiments are in excellent agreement with theory.

We have developed the theory for localizing two point targets, and are awaiting experimental results from the University of Rochester to validate that theory after which a journal article reporting our work on point-target localization will be prepared.
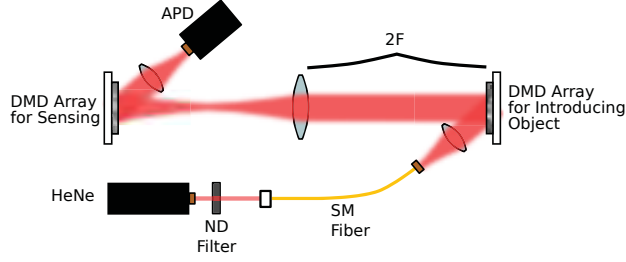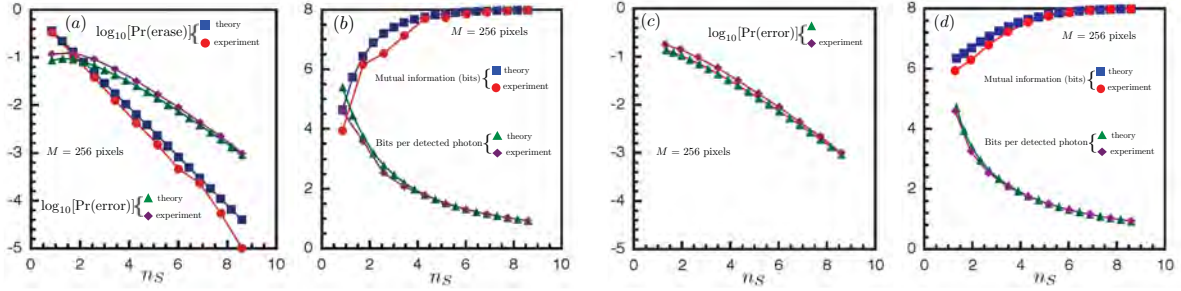
FIG. 27. Setup for point-target localization experiment.

FIG. 28. Pr(erase) and Pr(error) for dark-count limited operation [(a) and (c)] and MI and bpdp [(b) and (d)] without [(a) and (b)] and with [(c) and (d)] the PUD protocol.

D. Classical Ghost Imaging with Phase-Sensitive Light

The theory of partial coherence has a long and storied history in classical statistical optics. The vast majority of this work addresses fields that are statistically stationary in time, hence their complex envelopes only have phase-insensitive correlations. The quantum optics of squeezed-state generation, however, depends on nonlinear interactions producing baseband field operators with phase-insensitive and phase-sensitive correlations. Utilizing quantum light to enhance imaging has been a topic of considerable current interest, much of it involving biphotons, i.e., streams of entangled-photon pairs. Biphotons have been employed for quantum versions of optical coherence tomography, ghost imaging, holography, and lithography. However, their seemingly quantum features have been mimicked with classical-state light, questioning wherein lies the classical-quantum boundary. We have shown, for the case of Gaussian-state light, that this boundary is intimately connected to the theory of phase-sensitive partial coherence. In [119] we presented the theory of phase-sensitive partial coherence, contrasting it with the familiar case of phase-insensitive partial coherence, and used it to elucidate the classical-quantum boundary of ghost imaging. We showed, both theoretically [96] and experimentally [105], that classical phase-sensitive light produces ghost images most closely mimicking those obtained with biphotons. Sample results from our experiments are shown in Fig. 29. Consistent with theory for far-field ghost imaging, we found the phase-sensitive ghost image to be inverted, whereas the phase-insensitive ghost image was erect. We also observed similar spatial resolutions for both modes of operation, as predicted by their equal far-field coherence radii. The measured spatial resolution was roughly consistent with our observed ∼1 mm speckle radius. Note that a bright artifact at the center of the image, caused primarily by the sub-optimal 83% fill factor of our SLMs, prevented us from imaging effectively near that region.
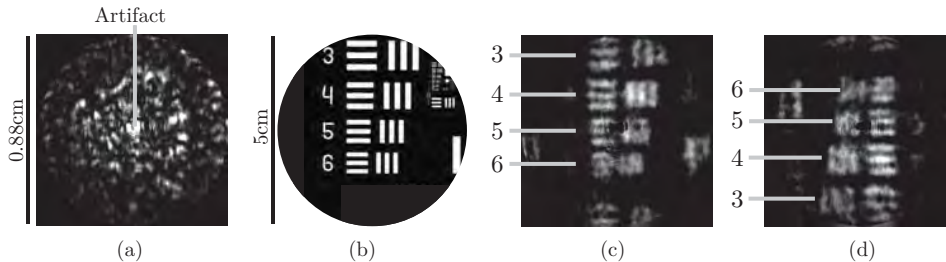
FIG. 29. (a) Sample single-frame speckle pattern as imaged by the CCD. The artifact is due to the SLMs' ∼83% pixel fill factor. (b) Portion of a USAF spatial-resolution transmission mask used as the object. (c) Phase-insensitive and (d) phase-sensitive far-field ghost images, each averaged over 18640 realizations. Background levels are clipped for improved visibility.

V. HIGH EFFICIENCY ORBITAL ANGULAR MOMENTUM IMAGING

Correlated optical sensing uses various types of correlations between pairs of photons or pairs of classical light beams to form an image or to detect specific object features. This includes techniques such as ghost imaging [120–123] and compressive ghost imaging [124, 125]. In ghost imaging, correlated light is sent through two different paths, one of which contains the target object and a bucket detector, and the other has a spatially-resolving detector but no object. Correlation between the outputs of each detector allows reconstruction of an image or sensing of particular features of the object [126].

In compressive sensing [127], the illumination of the target object is modulated by a number of known (usually random) spatially-varying transmission masks and the output is collected by a bucket detector. An estimation of the object or its features can be found by correlating the intensity of the detected light with the mask profile. Compressive ghost imaging uses the techniques of compressive sensing in a ghost imaging setup to attain more efficient imaging by means of post-processing, such as the recently proposed multiple-input ghost imaging [128], in which the bucket detector is replaced by a sparse-pixelated detector. The importance of compressive sensing stems from the fact that successful sensing of the object requires a number of measurements much smaller than the image size (measured in number of pixels), which has important implications for more efficient remote-sensing techniques. Here we wish to use correlated measurements in order to achieve something analogous to compressive sensing, but reducing the number of required measurements by a judicious choice of measurement instead of by post-processing. For complex objects, the method of this report may be combined with compressive sensing algorithms to achieve considerable reduction in the number of photons required to make optical measurements.

The use of orbital angular momentum (OAM) states in classical and quantum imaging techniques has been shown to provide additional effects that enhance the sensitivity to particular features of an object. For example, it has been shown that the use of OAM modes in phase imaging configurations increases edge contrast by using a spiral phase distribution as a filter [126]. In addition, digital spiral imaging (DSI) [129, 130] has been proposed as a technique in which the OAM basis is used to illuminate the object and to analyze the transmitted or reflected light. The two-dimensional spatial structure of the mode along with the high dimension of the OAM basis set leads to the probing of two-dimensional objects without obtaining a pixel by pixel reconstruction, analogous in some ways to the approach of compressive sensing.

In the following sections, we describe several related topics that are all based on combining the high dimensionality of the OAM based-sensing with the use of correlated two-photon states. In section VI, the first topic [131], entangled spiral imaging (ESI), is potentially suitable for OAM-based remote image reconstruction. The two-photon joint OAM spectrum provides object information concerning spatial symmetries, yielding more efficient object recognition than conventional imaging and sensing techniques. ESI can be viewed as similar to compressive sensing, in the sense that sampling is performed in randomly-chosen elements of a high-dimensional OAM basis, so that if the object has a compact representation in this basis, a small number of measurements will suffice to identify it. All of the methods described here are based on measurements of

the joint angular momentum spectrum of photon pairs. No measurements are made in position space.

Section VII covers a non-imaging variant of the same idea. The structure of the OAM states make them suitable for rapid detection of rotational symmetries, allowing objects of a know shape to be identified with few photons. In addition, this is useful for spotting objects that fail to have the desired symmetry; thus it may be useful for identifying faulty units on a production line or abnormal cells in a biological sample. Finally, in section VIII we use computer simulation to extend the main results from secs. VI and VII to the case of more complicated, asymmetric objects. We find that the informational capacity of our method remains above the one bit per photon limit.

A. Laguerre-Gauss modes

We decompose ingoing and outgoing beams in terms of optical Laguerre-Gauss (LG) modes. The LG wavefunction with OAM $l\hbar$ and with p radial nodes is [132]

$$u_{lp}(r, z, \phi) = \langle r, z, \phi | l, p \rangle \tag{21}$$

$$= \frac{C_p^{|l|}}{w(z)} \left(\frac{\sqrt{2}r}{w(z)} \right)^{|l|} e^{-r^2/w^2(r)} L_p^{|l|} \left(\frac{2r^2}{w^2(z)} \right)$$

$$\times \, e^{-ikr^2 z / \left(2(z^2 + z_R^2) \right)} e^{-i\phi l + i(2p + |l| + 1) \arctan(z/z_R)}, \tag{22}$$

with normalization $C_p^{|l|} = \sqrt{\frac{2p!}{\pi(p+|l|)!}}$ and beam radius $w(z) = w_0 \sqrt{1 + \frac{z}{z_R}}$ at z. $z_r = \frac{\pi w_0^2}{\lambda}$ is the Rayleigh range and the arctangent term is the Gouy phase. $|l, p\rangle$ and $|r, z, \phi\rangle$ are respectively the OAM eigenstate and the position eigenstate in cylindrical coordinates.

B. Digital spiral imaging

DSI [130] is a form of angular momentum spectroscopy in which properties of an object are reconstructed based on how it alters the OAM spectrum of light used to illuminate it (fig. 30). The input and output light may be expanded in LG functions, with the object acting by transforming the coefficients of the ingoing expansion into those of the outgoing expansion. Information about the transmission profiles of both phase and amplitude objects may be retrieved [130, 133].

The idea naturally arises of trying to use the measured OAM spectrum to reconstruct an image of the object. But, although a great deal of information may be obtained about the object in this manner, it is *not* sufficient to reconstruct a full image of the transmission or reflection profile. To see this, expand the output amplitude according to $\sum_{lp} A_{lp} u_{lp}$. Projecting out particular l and p values, the detector tells us the intensity of each component, allowing the $|A_{lp}|^2$ to be found, with no phase information retrieved. We thus have an incoherent imaging setup, with total detected intensity of the form $\sum_{lp} |A_{lp}|^2 |u_{lp}|^2$. But the quantities $|u_{lp}|^2$ are rotationally symmetric for all values of l and p (see the right-most panel of fig. 31). Any image built from them is also symmetric; variation of the object about the axis is lost. In contrast, the real and imaginary parts are *not* rotationally invariant (left two panels of fig. 31), so a coherent sum of the form $|\sum_{lp} A_{lp} u_{lp}|^2$ allows azimuthal structure to be reconstructed from the interference terms. For image reconstruction, we thus need to obtain a *coherent* superposition of amplitudes. This can be seen in fig. 32: an opaque square is placed in the beam and the two expansions (coherent and incoherent) are computed, assuming that only the $p = 0$ components are measured and keeping terms up to $|l_{max}| = 15$. In the left panel, where no phase information is assumed, the reconstructed image is rotationally invariant and there is no way to distinguish what the actual shape of the object was. In contrast, the coherent expansion on the right side of the figure produces a recognizably square output. *The phase information is vital in reconstructing the actual image shape.* A method for measuring the relative phases of the different OAM components and thereby reconstructing the image is discussed in section VI.

C. Entangled OAM pairs

The form of the two-photon state produced by spontaneous parametric downconversion (SPDC) is well-known. It is most often written as an expansion in the space of transverse linear momenta of the outgoing signal and idler:

$$|\Psi\rangle = \int d^2q_s \; d^2q_i \tilde{E}(q_s + q_i)\tilde{W}(q_s - q_i)\hat{a}^\dagger_{q_s}\hat{a}^\dagger_{q_i}|0\rangle. \tag{23}$$

Here, $\tilde{E}$ is the momentum space pump profile, and $\tilde{W}$ is the phase-matching function given by [134]

$$\tilde{W}(q_s - q_i) = \sqrt{\frac{2L}{\pi^2 k}}\text{sinc}\left(\frac{|q_s - q_i|^2 L}{4k}\right) \tag{24}$$
$$\times exp\left(-i\frac{|q_s - q_i|^2 L}{4k}\right),$$

where L is the thickness of the crystal and $k = \frac{\omega_p n_p}{c}$ is the magnitude of the pump momentum.

In our case, however, we wish to expand in the space of orbital angular momentum instead of transverse linear momentum. Consider a pump beam of spatial profile $E(r) = u_{l_0 p_0}(r)$ encountering a χ^2 nonlinear crystal, producing two outgoing beams via SPDC. For fixed beam waist, the range of OAM values produced by the crystal is roughly inversely proportional to the square root of the crystal thickness L [133]. We wish a broad OAM bandwidth, so we assume a thin crystal located at the beam waist ($z = 0$). The output is an entangled state [146], with a superposition of terms of form $u_{l'_1, p'_1} u_{l'_2, p'_2}$, angular momentum conservation requiring $l_0 = l'_1 + l'_2$. We will take the pump to have $l_0 = 0$, so that the OAM values just after the crystal are equal and opposite: $l'_1 = -l'_2 \equiv l$. The p'_1, p'_2 values are unconstrained, although the amplitudes drop rapidly with increasing p' values (see eq. (27) below). The output of the crystal may be expanded as a superposition of signal and idler LG states:

$$|\Psi\rangle = \sum_{l'_1, l'_2 = -\infty}^{\infty} \sum_{p'_1, p'_2 = 0}^{\infty} C^{l'_1, l'_2}_{p'_1 p'_2}|l'_1, p'_1; l'_2, p'_2\rangle\delta(l_0 - l'_1 - l'_2), \tag{25}$$

where the coupling coefficients are given by

$$C^{l'_1, l'_2}_{p'_1 p'_2} = \int d^2r \; \Phi(r)\left[u_{l'_1 p'_1}(r)u_{l'_2 p'_2}(r)\right]^*. \tag{26}$$
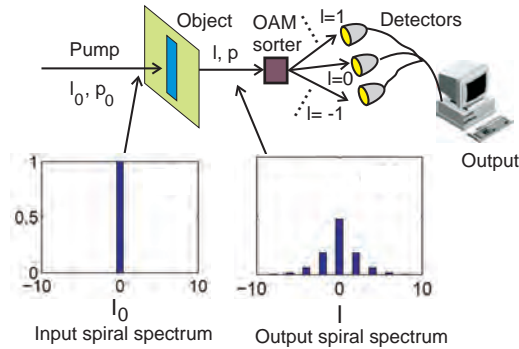

FIG. 30. *Digital spiral imaging: the presence of an object in the light beam alters the distribution of angular momentum values in the outgoing light.*
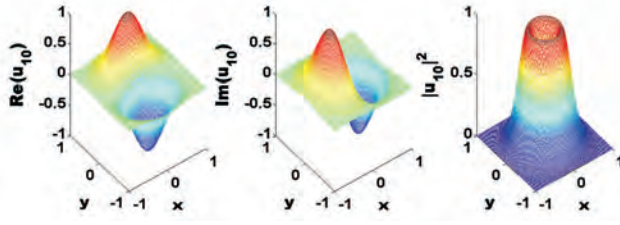
FIG. 31. *The real and imaginary parts of the Laguerre-Gauss function are not rotationally-invariant, in contrast to its absolute square. This is illustrated for the case of $l = 1$, $p = 0$, but is true generally.*
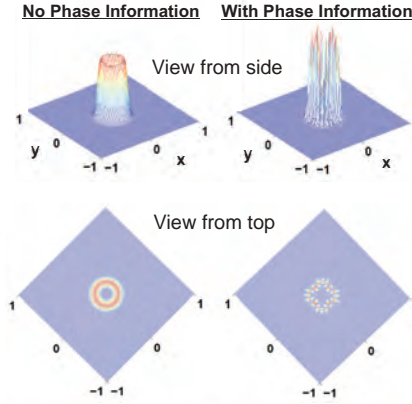

FIG. 32. *Incoherent (left) and coherent (right) expansions in Laguerre-Gauss functions, an opaque square object. In the former case, all variation of the object with angle around the axis is lost. ($p_{max} = 0$ and $l_{max} = 15$ assumed.)*

For the case of pump beam with $l_0 = p_0 = 0$ this gives the coefficients [133, 135]:

$$C_{p_1,p_2}^{l,-l} = \sum_{m=0}^{p_1} \sum_{n=0}^{p_2} \left(\frac{2}{3}\right)^{m+n+l} (-1)^{m+n} \tag{27}$$
$$\times \frac{\sqrt{p_1! p_2! (l+p_1)! (l+p_2)!} \, (l+m+n)!}{(p_1-m)!(p_2-n)!(l+m)!(l+n)! \, m! \, n!}.$$

VI. ENTANGLED SPIRAL IMAGING

A. Joint OAM spectra

We now investigate the use of two beams, rather than one, in combination with spiral imaging. The full benefits of doing this will emerge in section VI B. In the current section, we focus on examination of the OAM correlations. We begin with an entangled version, where the light source is parametric down conversion in a nonlinear crystal such as β-barium borate (BBO). Imagine an object in the signal beam (fig. 33). Since OAM exactly holds only in the paraxial case, we assume the signal and idler are produced in *collinear* down conversion, then directed into separate branches by a beam splitter. (Throughout the following we assume all beam splitters are 50-50.) Assume perfect detectors for simplicity (imperfect detectors can be accounted for by the method in [135]). For our purposes, either type I down conversion with a nonpolarizing beam splitter or type II with polarizing beam splitter will work, since the processes involved are not polarization-
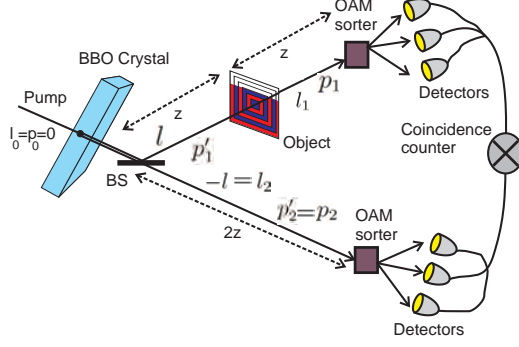
FIG. 33. *Setup for analyzing object via orbital angular momentum of entangled photon pairs.*

dependent. The only difference is that the use of type II with a polarizing beam splitter will increase the coincidence rate by a factor of two. (Note that both photons enter the beam splitter through the same port, rather than through opposite ports, so there is no possibility of complete destructive interference of the kind that leads to the HOM dip.)

Let $P(l_1, p_1; l_2, p_2)$ be the joint probability for detecting signal with quantum numbers l_1, p_1 and idler with values l_2, p_2. The marginal probabilities at the two detectors (probabilities for detection of a single photon, rather than for coincidence detection) are

$$P_s(l_1, p_1) = \sum_{l_2, p_2} P(l_1, p_1; l_2, p_2) \tag{28}$$

$$P_i(l_2, p_2) = \sum_{l_1, p_1} P(l_1, p_1; l_2, p_2). \tag{29}$$

Then the mutual information for the pair is

$$I(s, i) = \sum_{l_1, l_2 = l_{min}}^{l_{max}} \sum_{p_1, p_2 = 0}^{p_{max}} P(l_1, p_1; l_2, p_2) \tag{30}$$
$$\times \log_2 \left(\frac{P(l_1, p_1; l_2, p_2)}{P_s(l_1, p_1) P_i(l_2, p_2)} \right)$$

The most common experimental cases are when (i) the values of p_1 and p_2 are not measured (so all possible values of p_1 and p_2 must be summed, $p_{max} = \infty$), or (ii) only the $p_1 = p_2 = 0$ modes are detected ($p_{max} = 0$). Except when stated otherwise, we will use $l_{max} = -l_{min} = 10$ and $p_{max} = 0$.

If the transmission function for the object is $T(x)$, the coincidence probabilities $P(l_1, p_1; l_2, p_2) = |A_{p_1 p_2}^{l_1 l_2}|^2$ have amplitudes

$$A_{p_1 p_2}^{l_1 l_2} = C_0 \sum_{p_1'} C_{p_1' p_2}^{-l_2, l_2} a_{p_1' p_1}^{-l_2, l_1}(z), \tag{31}$$

$$a_{p_1' p_1}^{l_1' l_1}(z) = \int u_{l_1' p_1'}(x, z) \left[u_{l_1 p_1}(x, z) \right]^* T(x) d^2 x \tag{32}$$

where C_0 is a normalization constant. Here it is assumed that the total distance in each branch is $2z$ (see fig. 33).

That the object's size and shape affect the coincidence rate is easy to see. For example, fig. 34 shows the calculated spectrum when a single opaque strip of width d is placed in the beam. Fig. 35 shows the corresponding mutual information, assuming that only the $p_1 = p_2 = 0$ component is detected. In both figures, we see clear effects of changing an object parameter (the strip width). The central peak of the spectrum (fig. 34) broadens as d increases from zero, reducing the correlation between l_1 and l_2; the mutual
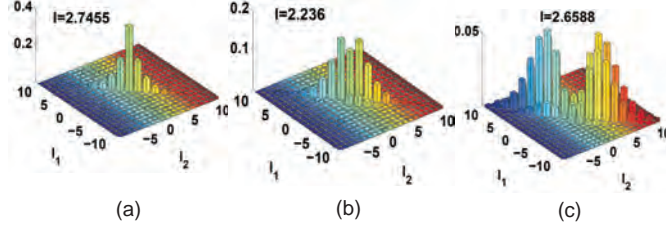
FIG. 34. *An opaque strip of width d placed in the signal path. The widths are (a) $d = .1w_0$, (b) $d = .9w_0$, (c) $d = 2.5w_0$. The outgoing joint angular momentum spectra are plotted. As the width increases, the peak in the spectrum broadens, then (at $d = w_0$) splits into two peaks.*
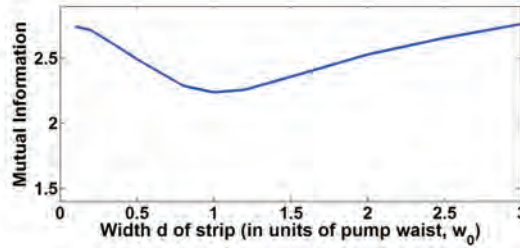

FIG. 35. *Mutual information versus width of opaque strip. The horizontal axis is in units of w_0. The minimum information occurs at $d = w_0$.*

information between them thus declines, as seen in the $d/w_0 < 1$ portion of fig. 35. But at $d/w_0 \approx 1$, the central peak in $\{l_1, l_2\}$ space bifurcates into two narrower peaks (right side of fig. 34); the information thus goes back up as the peaks separate, as indeed is the case in the $d/w_0 > 1$ region of fig. 35. If we continue to wider d, the two peaks once again broaden and the mutual information decays gradually to zero. In addition, the total intensity getting past the opaque strip will continue to drop, so coincidence counts decay rapidly.

B.   Imaging

The inability of DSI to produce images due to loss of phase information has been pointed out. Here we show that a variation on the ESI setup can be used to find the expansion coefficients *including phase*. Note that the phase of $A_{00}^{l_1 l_2}$ is the same as the phase of $a_{p'_1 0}$ for even p'_1 and differs from that of $a_{p'_1 0}$ by a factor of π for p'_1 odd [131]. Therefore, finding the phases of the coincidence detection amplitudes $A_{00}^{l_1 l_2}$ therefore suffices to determine the phases of all of the $a_{p'_1 0}$ coefficients.

The measurement of these phases is accomplished by inserting a beam splitter to mix the signal and idler beams before detection, as in fig. 36, erasing information about which photon followed which path. We then count singles rates in the two detection stages, rather than the coincidence rate. If value l is detected at a given detector it could have arrive by two different paths, so interference occurs between these two possibilities. The detection amplitudes in the two sets of detectors D_+ and D_- involve factors $A_+ \sim \left(1 + ia_{00}^{l_0 - l_2, l_1}\right)$ and $A_- \sim \left(i + a_{00}^{l_0 - l_2, l_1}\right)$, with detection rates $R_\pm \sim 1 + |a_{00}^{l_0 - l_2, l_1}|^2 \pm 2i \, Im \, a_{00}^{l_0 - l_2, l_1}$. From these counting rates, both the amplitudes and the relative phases of all coefficients can be found, allowing full image reconstruction.

Once the coefficients $a_{p'_1, p_1}^{l'_1, l_1}$ have been found from the coincidence rates, image reconstruction requires the inversion of eq. 32 to find the object transmission function $T(r)$. To facilitate this, we first define an operator
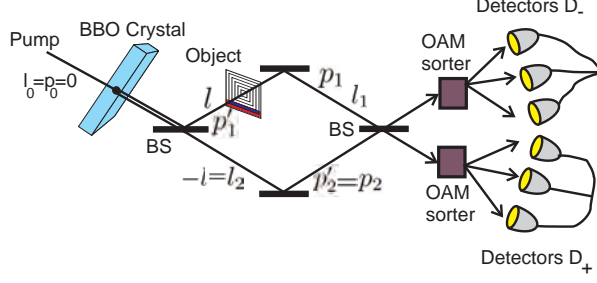
FIG. 36. *A configuration allowing image reconstruction via phase-sensitive measurement of entangled OAM content.*

$\hat{T}$ to represent the effect of the object on the beam. We may expand this operator in the position basis,

$$\hat{T} = \int d^2r \, d^2r' |r'\rangle T(r,r')\langle r| \tag{33}$$

$$= \int d^2r \, |r\rangle T(r)\langle r|, \tag{34}$$

where in the last line we assumed that the operator is a local operator, diagonal in the position space basis. The object function $T(r)$ in eq. 32 is then given by

$$T(r) = \langle r|\hat{T}|r\rangle. \tag{35}$$

Alternately, the object operator may be expanded in the Laguerre-Gauss basis,

$$\hat{T} = \sum_{ll'} \sum_{pp'} d^{l'l}_{p'p} \, |l'p'\rangle\langle lp|. \tag{36}$$

Making use of these definitions and of eq. 32, it follows immediately that

$$d^{l'_1,l_1}_{p'_1,p_1} = \langle l'_1 p'_1|\hat{T}|l_1 p_1\rangle = a^{l_1,l'_1}_{p_1,p'_1}. \tag{37}$$

(Note the reversal in the order of the indices.) Using this result in eq. 36, then applying eq. 35 and the fact that

$$u_{lp}(r) = \langle r|lp\rangle, \tag{38}$$

we find that determination of the $a^{l'_1,l_1}_{p'_1,p_1}$ coefficients is equivalent to reconstructing the object, since

$$T(r) = \langle r|\hat{T}|r\rangle = \sum_{ll'} \sum_{pp'} a^{l'_1,l_1}_{p'_1,p_1} u_{l_1 p_1}(r) \left[u_{l'_1 p'_1}(r)\right]^* \tag{39}$$

An example is shown in fig. 37, in which the object composed of a single opaque band (strip width equal to .5 times beam waist) is reconstructed using eq. 39, for different values of l_{max} and p_{max}. We see that as the number of values of l included is increased, a valley appears in the transmission profile at the location of the opaque band, and gradually becomes sharper and more pronounced as we increase l_{max}. Note that the number of measurements required for the image reconstruction (one for each l and p combination) is much smaller than is required by standard methods (one measurement per pixel).
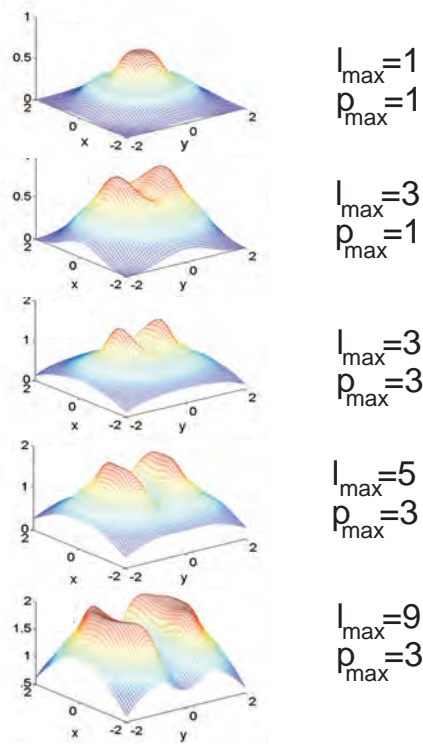
FIG. 37. *Reconstruction of the transmission profile of an object with a single opaque band of width .5w₀ for several values of l_{max} and p_{max}. We see the shape of the object appear and begin to sharpen as more l and p values are included.*

C. Correlated Spiral Imaging (CSI)

In recent years, it has been shown that ghost imaging and other "quantum" two-photon effects may be carried out using classically-correlated sources [123, 136–140]. It is apparent that the same is true in the case of correlated spiral imaging: classical OAM correlation, rather than entanglement, is sufficient. The essential point in the present case is having two spatially separated beams such that if the OAM detected in one beam is known, then the OAM reaching the object can be predicted. So all that is needed is strong classical correlation or anti-correlation between the OAM in the two arms.

The classical analog of apparatus of fig. 36 is shown in fig. 38. At the left, the system is illuminated with light that has a broad range of OAM values (a broad spiral spectrum). The beam is split, with one copy passing through the object, and the other entering the reference branch. The two beams are mixed at the beam splitter, then the OAM content at the two detectors is measured. The coefficients $C_{p_1,p_2}^{l_1,l_2}$ will no longer be given by eq. 27, but instead will have values determined by the properties of the specific input beam being used. The mutual information between the classical beams may be defined just as in eq. 30.

The classical configuration of CSI has a number of practical advantages over the entangled version: alignment issues are greatly reduced, single photon detectors are not needed, and much higher brightness and counting rate may be obtained. There is one problem that arises, however, which is not present in the entangled case: if a broad spiral spectrum is used for the illumination, then there is no intrinsic correlation between the OAM value l_2 in the reference branch and the value l_1' that occurs between the source and object. Without this correlation, the value of l_1' is unknown and so the change in l produced by the object is also unknown. On the other hand, instead of a broad spiral spectrum, we may send in single OAM values, one at a time, building up the OAM correlation function one value of l_1' at a time. But this slows the process of image reconstruction considerably: a range of OAM values needs to be scanned over, one after another,
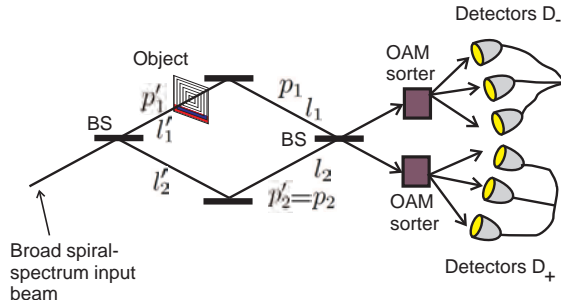
FIG. 38. *A classical version of correlated spiral imaging. An input beam with a broad range of OAM values is split at a beam splitter, sending a portion through the reference branch, and the rest to the object.*

changing a spiral phase plate or some other type of OAM filter multiple times in each run. In a kind of quantum parallelism, the entangled version can send in a broad range of values simultaneously, and the entangled nature of the source will automatically ensure that the pairs detected are of opposite initial OAM if a short enough coincidence time window is used. In any case, whether the classical or entangled version is used, two correlated beams are necessary in order to reconstruct the relative phases of the various OAM amplitudes.

D. Compressive imaging

Recent years have shown an explosion of interest in compressive sensing [127, 141–144], including compressive ghost imaging [124]. The basic idea is that most images are very sparse when expanded in an appropriate basis, with the vast majority of expansion coefficients being very small. So if a sampling procedure is used that only measures the relatively small number of large expansion coefficients and neglects the rest, the image may be reconstructed from a very small number of measurements, often much smaller than naively expected from the Shannon-Nyquist theorem.

For simple objects with well-defined rotational symmetries, the image can be reconstructed with fewer photons than with pixel-by-pixel imaging, in rough analogy to compressive sensing. For more complicated objects, correlated OAM imaging may be combined with compressive sensing for further increases in efficiency.

VII. EXPERIMENTAL DEMONSTRATION USING CORRELATED OAM SPECTRUM FOR OBJECT IDENTIFICATION

Now, rather than attempt to reconstruct an image, in the remainder of this report we focus on studying the properties of the joint OAM spectrum in the presence of the object, and then use knowledge of these properties to identify objects. The goal is to make measurement entirely in the OAM basis in order to identify objects, without attempting image reconstruction. Significant information from the object is available by measuring the *joint* OAM spectrum of two-photon states, giving rise to novel features outside the main OAM-conserving diagonal [146], due to the interaction with the object. In order to use this information for object identification, the relationship of the joint OAM spectrum to the azimuthal Fourier coefficients is derived. In the field of remote sensing this capability enables efficient object identification without the requirement of a full two-dimensional image. The present demonstration also represents the first correlated OAM-based sensing experiment representative of a practical remote-sensing setup, in which physical objects completely detached from any optical component are used as targets for identification. All previous experiments on correlated sensing using OAM states ([126] for example) have used simulated objects drawn on a spatial light modulator (SLM). High sensitivity to the spatial properties of the targets is obtained, due to the high amount of information per photon attainable by the use of OAM states[145].
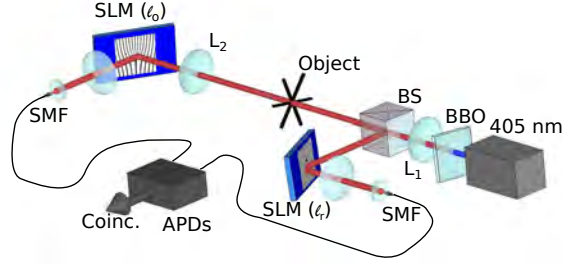
FIG. 39. *Setup for the sensing of the object via orbital angular momentum of correlated photons.*

A. Experiment

The experimental setup schematic is shown in Fig. 39. A 1.5 mm-thick BBO crystal is pumped with a 30 mW diode laser at 405 nm. To measure the OAM spectrum, an SLM is used to display computer generated holograms. Discrimination between individual radial modes $p \neq 0$ is not easily accomplished by current experimental techniques [147], therefore only the case $p = 0$ will be considered for simplicity and the p index will be dropped hereafter. The holograms change the winding number l of LG light, while two lenses and a mono-mode fiber selectively couple modes with $l = 0$ [135]. One condition for obtaining $a_{kl} = \langle k|\hat{T}|l\rangle$ without any post-processing is that the crystal, object, and detector planes are imaged onto one another, in order for all expansions to be carried out onto LG basis states of the same waist size. In the current setup, which is a representative demonstration of remote sensing, lens L2 images an object at a fixed distance onto the plane of the SLM. In a general remote sensing application, L2 will be placed appropriately so that all targets will be at distances greater than the hyperfocal length of L2. Therefore, all objects will be imaged in focus onto the SLM plane.

Fig. 40 (a–b) illustrates the measurement of the natural SPDC OAM joint spectrum C_l. The sign of l_r is flipped to compensate for the odd number of mirrors in order to reflect the OAM conservation in SPDC. In the absence of an object, only the main diagonal terms which match the total OAM content of the pump appear, without any non-diagonal terms observed [148]. Experimental data represent the mean of four measurements, from which the standard deviation is calculated and displayed as error bars. Fig. 40 (c) shows the histogram of the diagonal and the cross section for $P(l_o|l_r = 0)$.

States with $l_r + l_o \neq l_p$ can have non-zero probability. When there was no object present, these states were forbidden by OAM conservation [133, 135, 146–148]; they represent the interaction of the correlated OAM state from SPDC with the object. These new non-conservation elements have a contribution to its total OAM content from the object $m = l_o + l_r - l_p$, and carry direct information of the m-fold rotational symmetries of the object at a $|l_r| + |l_o|$ radius. For instance, considering a Gaussian pump ($l_p = 0$), the measured joint spectrum with a m-fold rotationally symmetric target will show strong components with total OAM content $\pm m$, and higher harmonics with lower amplitudes at nm, $|n| > 1$ being n an integer.

To demonstrate the capability and high sensitivity of this new technique for object recognition, the joint spectra using targets with different rotational symmetries are analyzed. An object can impart extra features in the joint spectrum according to its azimuthal Fourier series at different radii. In particular, an object with four-fold rotational symmetry will have strong signatures corresponding to a total OAM content of ± 4. To fit the scale of our setup, opaque strips 175 μm ($0.83w_0$) thick, are placed in the object arm arranged with specific rotational symmetries. These targets behave as transmission masks that block light, and the integration time is adjusted in order to have a similar amount of counts with respect to the case of no object. Test objects with well-defined dominant four- and six-fold symmetries shown in Figs. 41 (a) and 42 (a) considerably modify the joint spectra [Figs. 41 (b) and 42 (b)] by adding extra diagonals with total OAM content of $l = \pm 4$ and $l = \pm 6$, respectively. The two objects are clearly distinguished by the unique features of their respective joint spectrum.
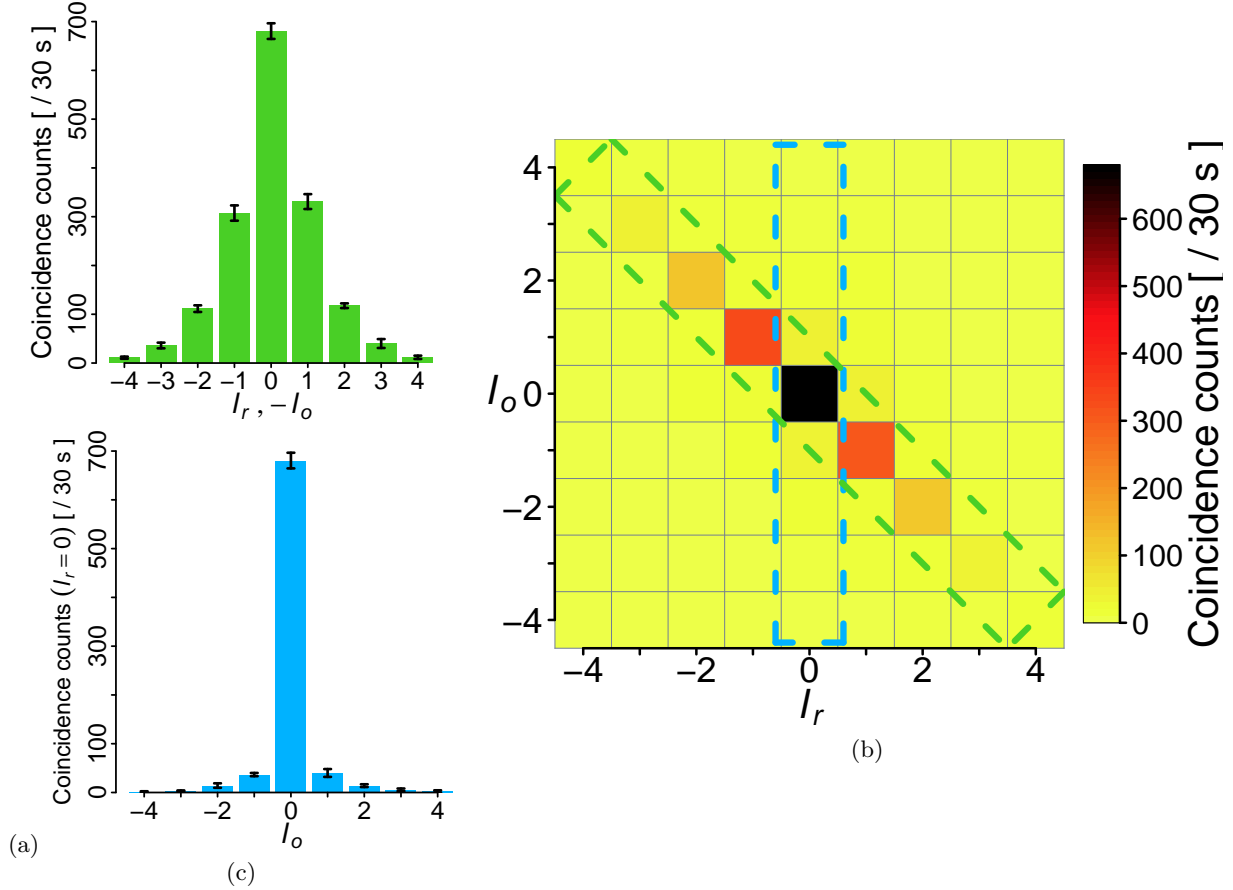
FIG. 40. (a) Histogram of the main diagonal and (b) complete two-dimensional OAM joint spectrum for our config-
uration with no object. (c) Corresponding $P(l_o|l_r = 0)$ cross section. The blue (vertical) box in the joint spectrum
denotes the section in (a) and the green (diagonal) box denotes the section in (c).

B. Discussion

More insight about the structure of each object is acquired by analyzing specific non-diagonal cross sections
such as the two in Figs. 41 (d) and 42 (d). Previous works regarding the OAM joint spectrum of SPDC
[131, 133, 135, 147, 148] have considered only the main diagonal elements corresponding to the conservation
of OAM with respect to the pump [Figs. 41 (c) and 42 (c)]. In the current case, there is an interaction
between the OAM from SPDC and the spatial features of the object resulting in contributions which require
analysis of the complete two-dimensional joint spectrum.

The non-diagonal cross section $P(l_o|l_r = 0)$ shown in Fig. 41 (d) indicates significant four-fold symmetry as
the dominant feature of the object. This feature can be uniquely attributed to the object because it is outside
the diagonal that corresponds to the conservation of OAM in SPDC. Fig. 42 (d), showing the same non-
diagonal cross section for the six-fold symmetric object, exhibits richer features associated with the higher
complexity of the object. Specifically, apart from the dominant six-fold symmetric contribution, a three-fold
contribution is observed. At the center of the object [Fig. 42 (a)], the three stripes are displaced with respect
to the center. This forms a triangular shape in a small region with three-fold rotational symmetry. This
small deviation from a strict six-fold symmetry is readily observed in the non-diagonal cross section by the
appearance of significant contributions at $l_o = \pm 3$.

Although the cross section $P(l_o|l_r = 0)$ is a good indicator of the presence of different symmetries, it is
necessary to analyze the full two-dimensional OAM joint spectrum to extract the relative contribution of
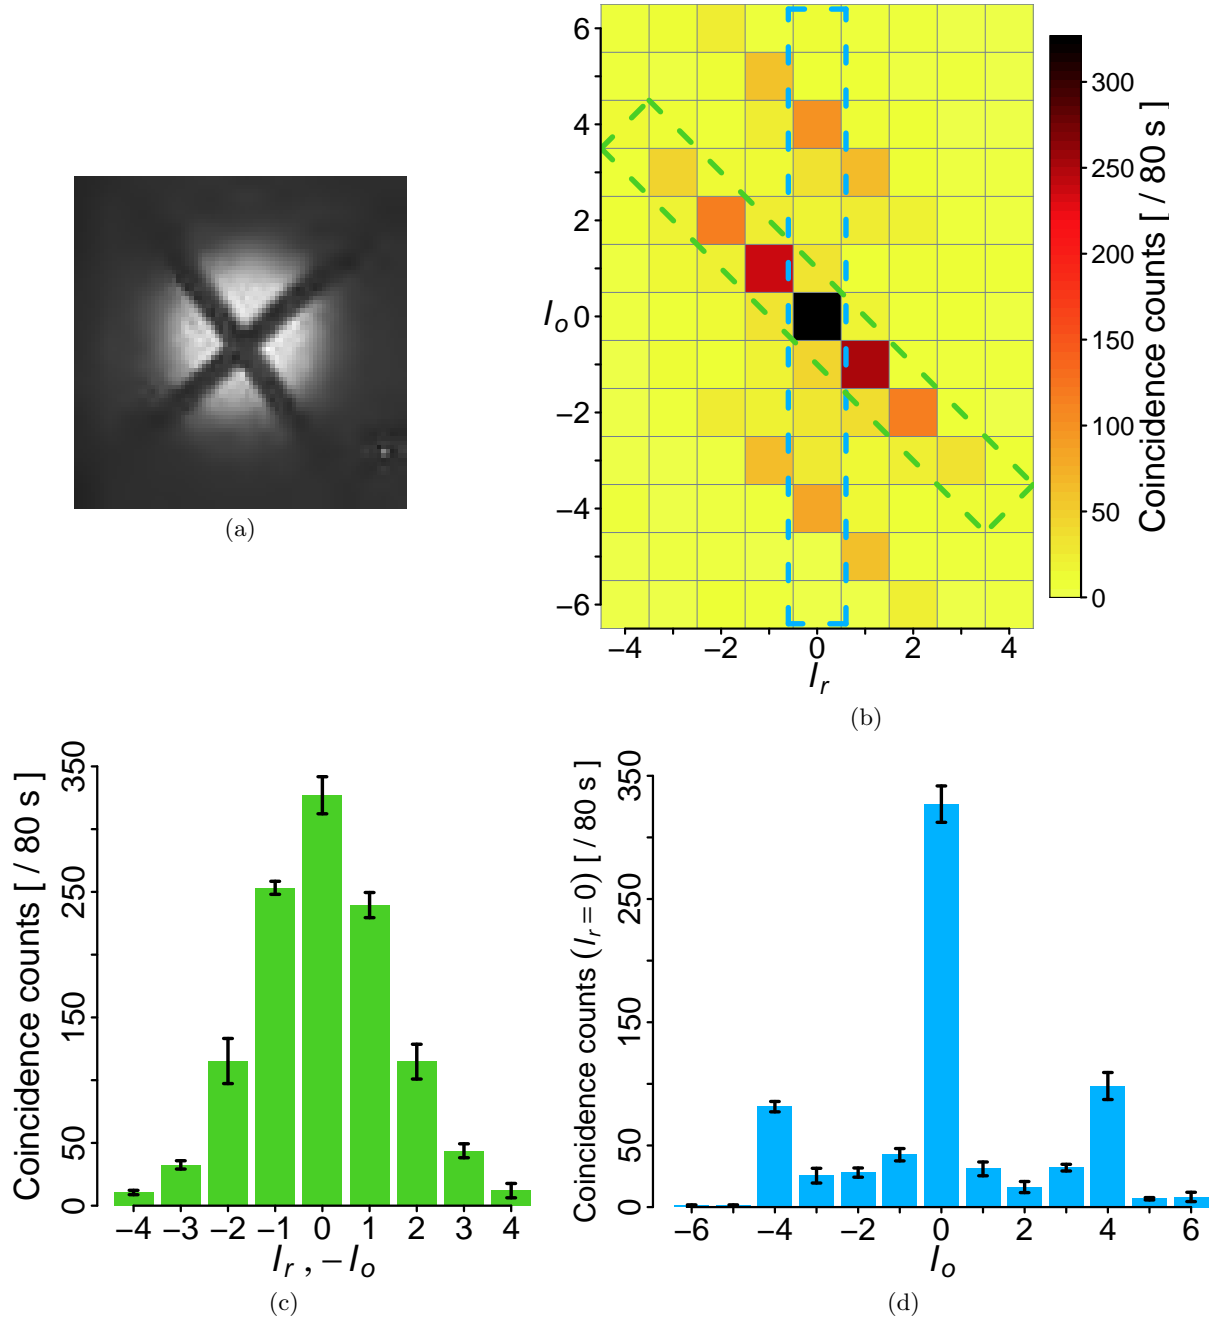
FIG. 41. (a) Image of the cross used as target, (b) the experimental joint spectrum, (c) a histogram of the joint spectrum main diagonal and (d) $P(l_o|l_r = 0)$ cross section of the joint spectrum.

each symmetry component. For example, Fig. 42 (d) indicates a relative contribution of six- and three-fold symmetries at a ratio of approximately 2:1. However, to judge the relative size of the areas within the object that exhibit each symmetry, the total contribution of the entire diagonals should be considered. In this case, the joint spectrum in Fig. 42 (b) clearly shows that the region with three-fold symmetry is significantly smaller than the region with six-fold symmetry.

It should be pointed out that rotating the object about the azimuthal axis simply shifts the phase of the
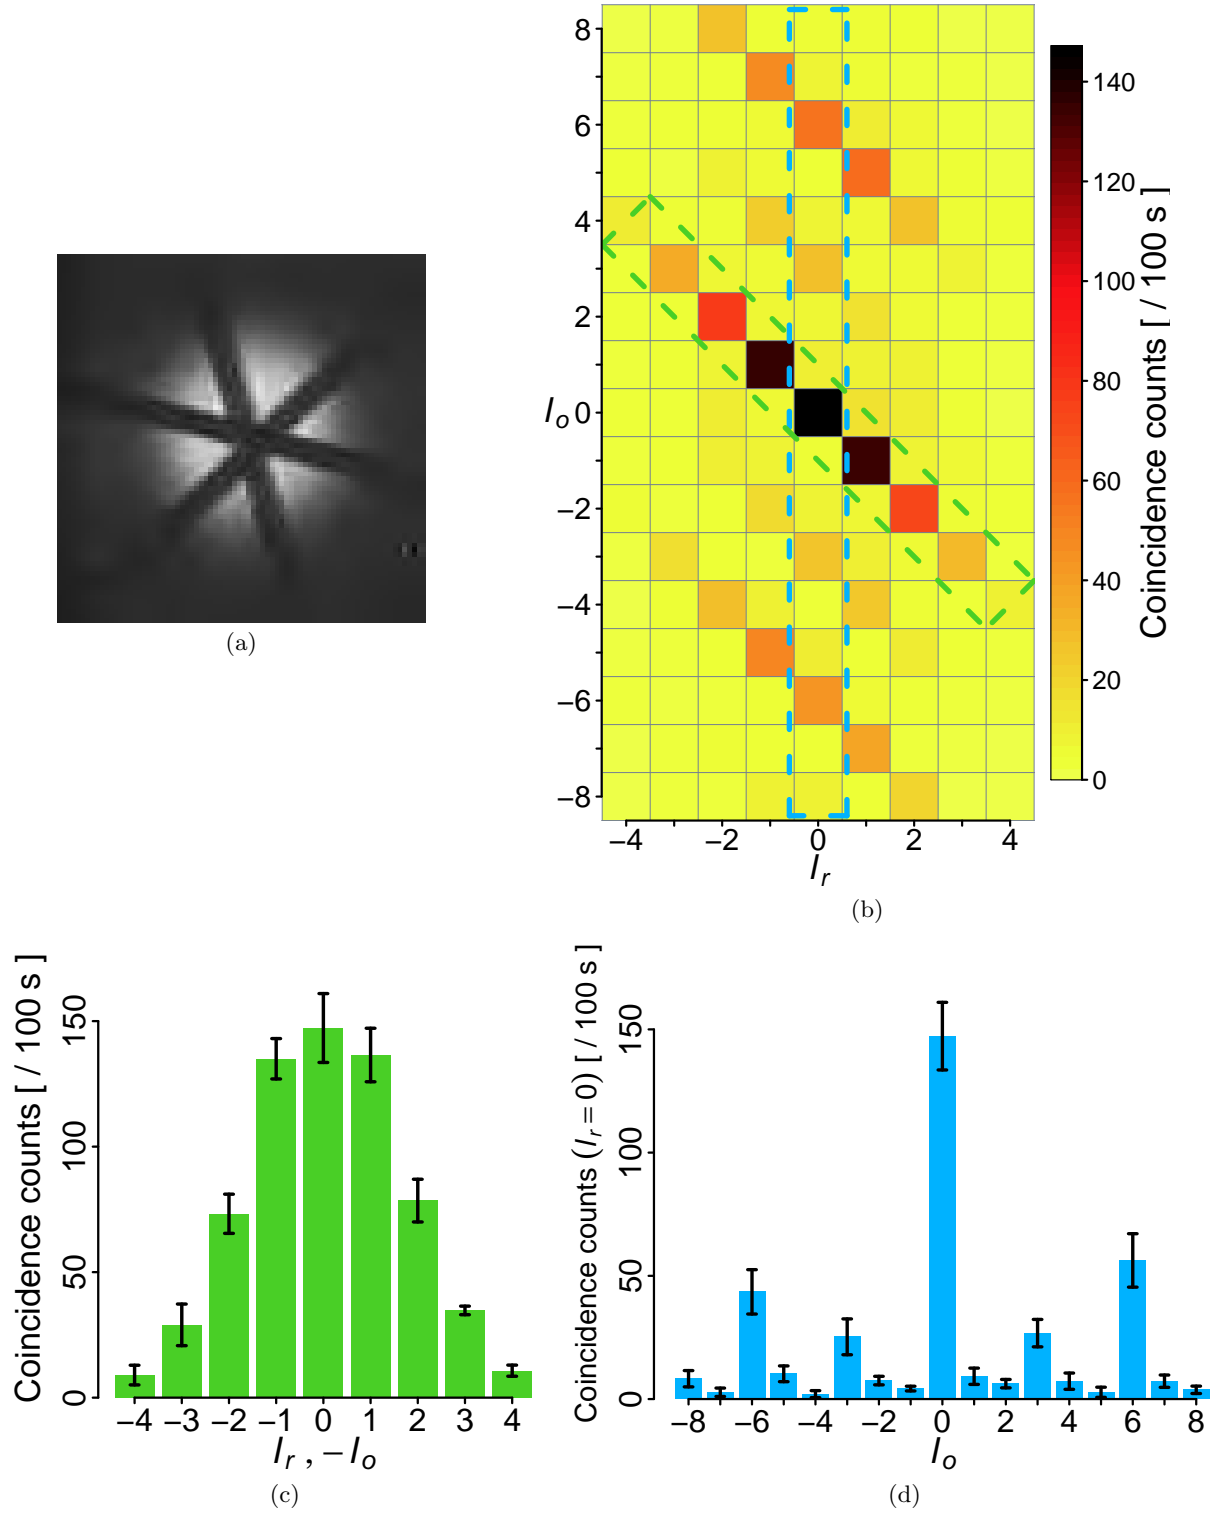
FIG. 42. (a) Image of the three-arm cross used as target, (b) the experimental joint spectrum, (c) a histogram of the joint spectrum main diagonal and (d) $P(l_o|l_r = 0)$ cross section of the joint spectrum.

field by a constant, which will have no effect on the correlation rate. Therefore, if an object is rotated, the method will correctly continue to identify it as the same object.

These results demonstrate that an object imprints its own characteristic features onto the joint OAM joint spectrum of SPDC. New elements arise that do not fulfill the OAM conservation condition required by SPDC alone, but which are affected by the object as well. The capability of ESI for object recognition is demonstrated by these results including high sensitivity to small symmetry components. This represents a realistic remote sensing application using physical objects detached from any optical components. Although the experiment was conducted with a transmissive object, a trivial modification of the optical setup (simply changing the position of the lens, SLM, and object arm detector) makes it suitable for remote sensing of reflective targets at arbitrarily large distances.

VIII. SYMMETRY, INFORMATION, AND IMAGING OF COMPLEX OBJECTS

Thus far, all objects considered have been exactly centered with respect to the beam axis, and have exhibited simple rotational symmetries. It has been shown that the non-zero off-diagonal elements of the joint OAM coincidence spectrum clearly indicate the presence of N-fold rotational symmetries in the object (see fig. 42), and that the setup in fig. 36 can be used to reconstruct the object image by recovering the relative phases of the coincidence amplitudes (see sec. VI B). In this section we implement numerical simulations of the setup described by fig. 36. Recall that the setup described there requires measurement of both coincidences and singles rates in order to recover the magnitudes and phases of the amplitudes in eq. 32, which are then used as the expansion coefficients that describe the object in eq. 36. In the simulations below, we use digitized representations of opaque objects to directly compute eq. 32 in order to study the effects – on spectral signature, image reconstruction accuracy, and mutual information – of translating the target objects off-axis with respect to the beam's center. Further, in addition to objects with simple symmetries, we consider more complicated geometries [149].

A. Mutual information and symmetry

Fig. 43 shows the computed mutual information for several simple shapes. It can be seen that I depends strongly on the size and shape of the object, so that for object identification from among a small set a comparison of the I values rather than of the full probability distribution may suffice.

If the object has rotational symmetry about the pump axis, then its transmission function $T(r)$ depends only on radial distance r, not on azimuthal angle ϕ. Thus in the case of rotational symmetry, the second arm becomes irrelevant from an information standpoint, since the mutual information simply reduces to the Shannon information extracted from one arm alone. In this sense, the quantity $\mu(L_1, L_2) \equiv |I(L_1, L_2) - S_1(L_1)|$ is an order parameter, capable of detecting breaking of rotational symmetry.

More generally, suppose that the object has a rotational symmetry group of order N; i.e., it is invariant under $\phi \to \phi + \frac{2\pi}{N}$. From eqs. (22) and (32) it follows that the coefficients must then satisfy $a_{p_1' p_1}^{l_1' l_1} = e^{\frac{2\pi i}{N}(l_1' - l_1)} a_{p_1' p_1}^{l_1' l_1}$, which implies $a_{p_1' p_1}^{l_1' l_1} = 0$ except when $\frac{l_1' - l_1}{N}$ is integer. When N goes up (enlarged symmetry group), the number of nonzero $a_{p_1' p_1}^{l_1' l_1}$ goes down; with the probability concentrated in a smaller number of configurations, correlations increase and mutual information goes up. This may be seen in the three right-most objects of fig. 43, for example.

B. Spectral Signature and Imaging of Complex Objects

The experimental results discussed in sec. VII B indicate the role of symmetry in an object's joint OAM spectrum. It is worth noting that the objects used there, while having width much smaller than the beam waist, had length that extended far beyond the beam radius at any point. Fig. 44(b) shows the joint spectrum of a simple 5-pointed opaque star (with 5-fold rotational symmetry) whose dimensions are confined entirely
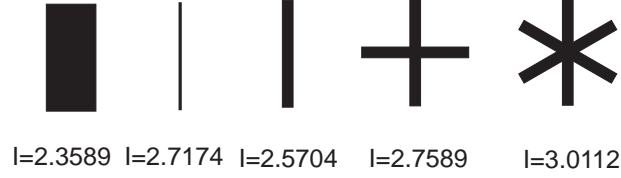
I=2.3589 I=2.7174 I=2.5704 I=2.7589 I=3.0112

FIG. 43. *The mutual information depends strongly on size and shape of the object. Here, the two objects on the left have widths* $1.5w_0$ *and* $.2w_0$*; all other widths are* $.4w_0$*.*
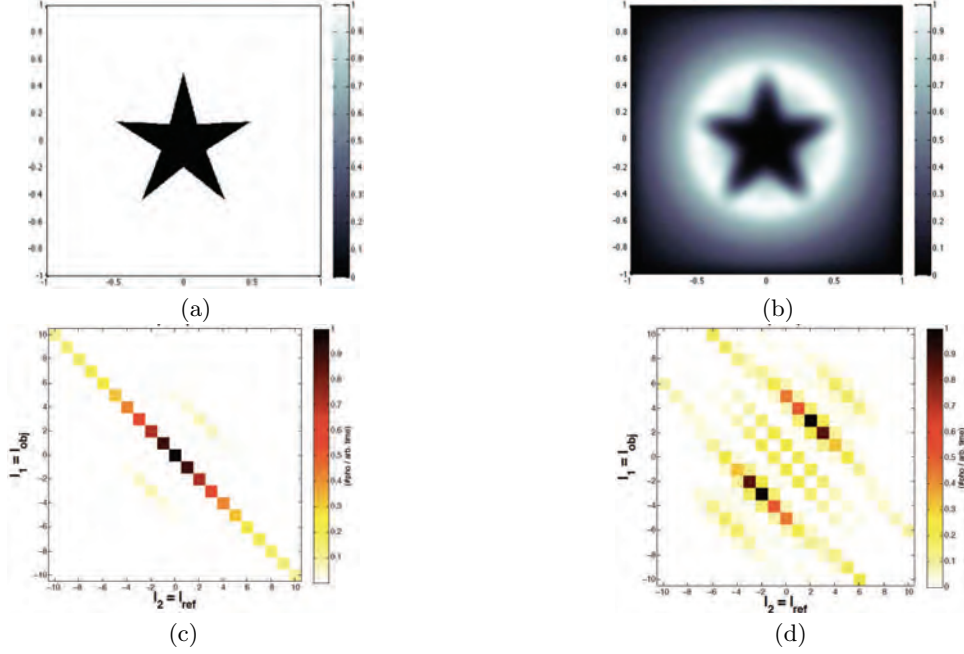

FIG. 44. (a) Opaque star object of max width $0.9w_0$ and, (b) the ESI reconstruction using $l_{max} = 10$, $p_{max} = 7$; (c) The joint OAM spectrum of the star, having summed over all p, and (d) the same spectrum with the conservation diagonal removed.

within the beam. The object's lack of radial extension causes a decrease in magnitude of the $l = \pm 5$ components of the joint spectrum, since the LG modes of higher momentum (higher l) do not interact with the object . Consequently, the object's spectral signature in the off-diagonal comments of the joint OAM spectrum becomes visually less obvious. However, as fig. 44(d) shows, by setting the diagonal components of the joint spectrum to 0 and rescaling the colormap used to view the spectrum, the off-diagonal contributions become much more visible. Since it is these off-diagonal contributions that carry the extra information upon which the ESI setup is based, in order to amplify the contribution of off-diagonal spectral components, we will zero out the conservation diagonal (states with $l_o = -l_r$) for the remaining object spectra simulated in this report. The image reconstructions will include the contributions of the $l_o = -l_r$ states.

Figs. 45(b) and 46(b) simulate the ability of the ESI method to image objects with more complicated, less symmetric transmission functions $T(x)$ than previously considered. Successful imaging of complicated objects requires the detectability of states with $p > 0$, since the expansion basis for $\hat{T}$ depends upon distinct contributions from each (p', p) combination in the set of basis states (see eq. 36). The joint spectra shown in figs. 46(c) and 45(c) are clearly less compact than those of simpler objects, like the star. This is to be expected when one views complicated objects as a superposition of many simpler, symmetric objects translated with respect to the beam axis: as shown below (sec. VIII C), translation with respect to the beam
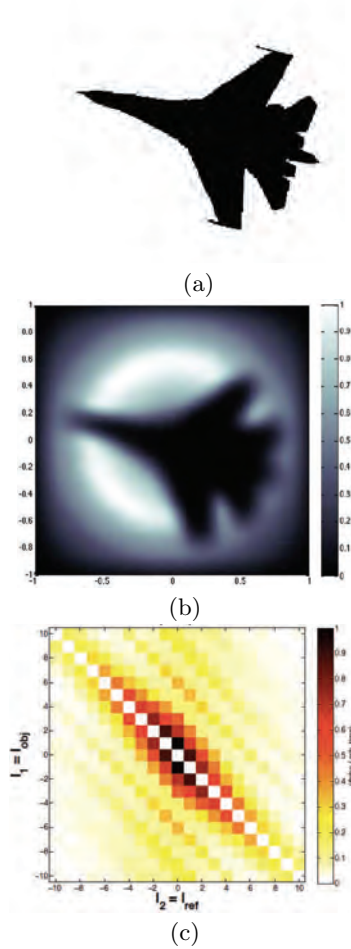
(a)

(b)

(c)

FIG. 45. (a) Opaque fighter jet object of max width w_0 and, (b) the ESI reconstruction using $l_{max} = 10$, $p_{max} = 7$; (c) The joint OAM spectrum of the fighter jet, summed over all p with the conservation diagonal removed.

axis, even for simple objects, spreads the joint spectrum.

C. Off-Axis Translation and Mutual Information

In figs. 47, 48, and 49 we show the image reconstruction and spectral signatures of the same objects shown in 44, 45, and 46 respectively *after having been shrunk by a factor of 4 and translated radially with respect to the beam axis* by approximately $0.7w_0 - 0.9w_0$. The effect of translation is most obvious in the case of the star, whose centered spectral signature is quite sparse compared to those of the tank or fighter jet. Namely, we observe that translation with respect to the beam axis causes a *spreading* in the spectral distribution. Although the exact dynamics of the spectral spread caused by translation vary from object to object, we note that once the object is sufficiently far from the beam center – not surprisingly – the conservation diagonal is all that remains, all off-diagonal components going to zero.

Given the variation in spectral signature as the object is translated through the beam field, we expect to see a corresponding variation in the mutual information carried by the components of the joint OAM spectrum. To calculate this change, we simulated the spectra of the above objects several times, linearly translating them with each iteration, starting from the beam center and ending effectively outside of the
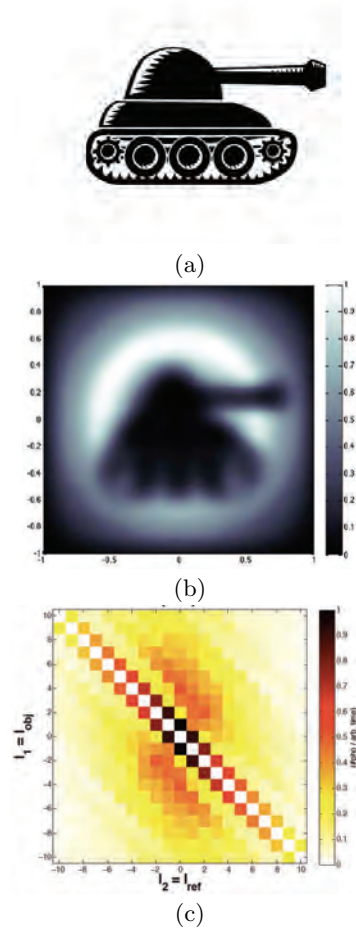
(a)

(b)

(c)

FIG. 46. (a) Opaque tank object of max width $0.7w_0$ and, (b) the ESI reconstruction using $l_{max} = 10$, $p_{max} = 7$; (c) The joint OAM spectrum of the tank, summed over all p with the conservation diagonal removed.

(a)   (b)
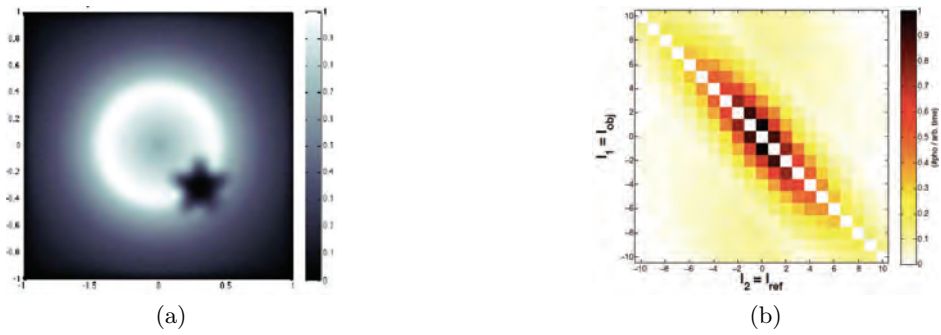
FIG. 47. (a) The ESI reconstruction of a translated opaque star, using $l_{max} = 10$, $p_{max} = 7$; (b) The joint OAM spectrum of the translated star, summed over all p with the conservation diagonal removed.

beam field completely. For each position the mutual information was calculated using eqn. 30, and the results are plotted as a function of distance in fig. 50. Since we are in this study primarily interested in the information content of the heretofore unconsidered off-diagonal components of the joint OAM spectrum, we

FIG. 48. (a) The ESI reconstruction of a translated opaque fighter jet, using $l_{max} = 10$, $p_{max} = 7$; (b) The joint OAM spectrum of the fighter jet, summed over all p with the conservation diagonal removed.
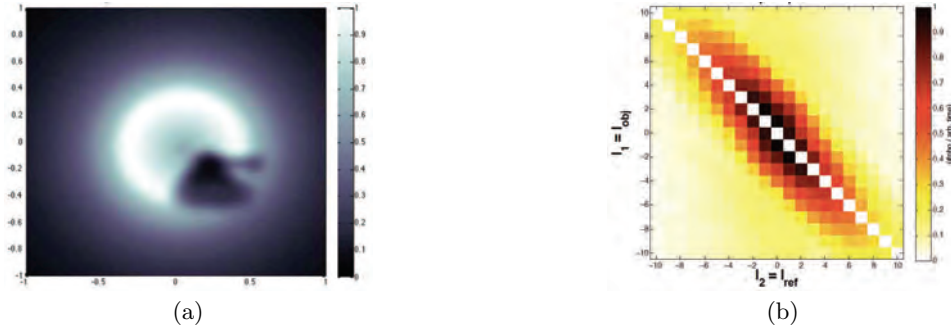
FIG. 49. (a) The ESI reconstruction of a translated opaque tank, using $l_{max} = 10$, $p_{max} = 7$; (b) The joint OAM spectrum of the tank, summed over all p with the conservation diagonal removed.
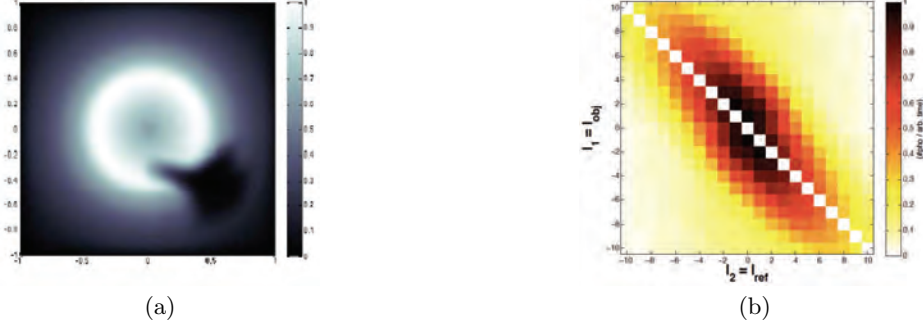
again zero out the conservation diagonal for all spectra so that *the mutual information calculated represents information carried exclusively by off-diagonal components of the spectrum.*

We see that *even for complex objects near the beam center, the mutual information carried by off-diagonal*

FIG. 50. *Mutual information carried by off-diagonal components of joint OAM spectrum, for various objects, as a function of distance from beam center with $l_{max} = 10, p_{max} = 5$; increasing p_{max} will increase the mutual information substantially. Note that each object's off-diagonal information content exceeds one bit per photon at the beam center.*

components of the joint OAM spectrum exceeds one bit per photon. As expected, the information goes to zero as the object moves sufficiently far from the b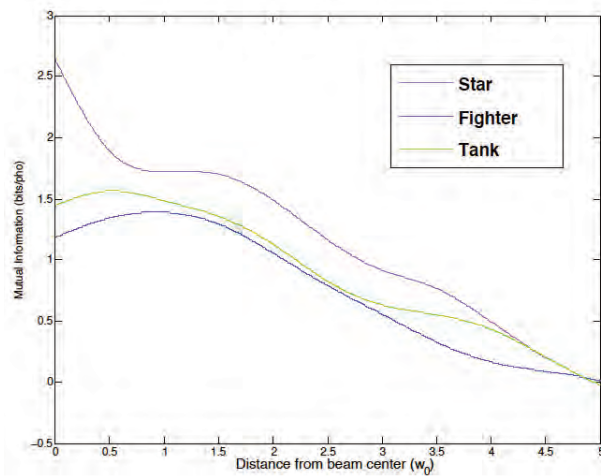eam center. Note that the simpler the object here – the star – carries the most off-diagonal information, consistent with the argument made in sec. VIII A, that enlarged symmetry groups cause an increase in correlations which in turn causes the mutual information to go up. In fact, as we increase p_{max} and the objects' symmetries are better approximated, the mutual information for each object goes up. In fig. 50, $p \in (0, 3)$ with the star's $I_{max} \approx 2.6$ bits/pho. Increasing p_{max} to 7 gives an $I_{max} \approx 3.3$ bits/pho in off-diagonal components.

D.   Discussion

The above simulations demonstrate the informational capacity of off-diagonal components in the joint OAM spectra. We have exploited this capacity for the purposes of imaging and object identification by way of the joint OAM spectral signature. Current experimental barriers, namely the inability to easily detect $p > 0$ modes at the single photon level, present difficulties in physically implementing the experimental apparatus required to recover the phases of the amplitudes needed for image reconstruction. However, as our simulations indicate, such an apparatus would be capable of using the information contained in the off-diagonal components of the joint OAM spectrum to remote image unknown objects *without any record of the spatial distribution of the photons measured.*

IX.   CONCLUSIONS

The experimental results reported above, along with the spectral signatures simulated in the final section of this report, rely only on coincidence measurements. This means that, where a set of objects with unique signatures or symmetries is in question, our method can be used to detect the presence or absence of objects in question in relatively few measurements as compared to pixel-by-pixel imaging methods.

A number of novel applications suggest themselves based on the results above. For example, note that if the object is rotated, the outgoing OAM states simply pick up an overall phase that does not affect the joint OAM spectrum. This could be useful, because it allows a rapidly rotating object to be imaged from its OAM spectrum using slow detectors. In some circumstances, this may be be less expensive and more practical than the use of high speed cameras.

The high mutual information capacity of off-diagonal OAM spectral components also makes our method well suited for sensing rotational symmetries in few measurements. Due to the fragility of OAM states, the advantages of our setup may best be exploited in small scale biological or production contexts. For example, the scanning of a biological sample using correlated OAM measurements may enable efficient detection of the presence or absence of certain structures based on the comparison of theoretical and observed coincidence rates of off-diagonal spectral components. And, since objects sufficiently far from the beam center do not affect the coincidence rates, as seen by the mutual information plots in 50, we can be confident that a sufficiently small beam waist will yield accurate spectra. Biological *apoptosis* (so-called programmed cell death) is one context in which the presence or absence of cell symmetries plays an important role, since apoptotic cells lose their symmetry, and an abundance of such cells may indicate a cancerous sample.

Before the research reported above, off-diagonal components of joint OAM spectra were not considered from an informational standpoint. We have found that not only do these components carry information, but they do so at rates which can well exceed the bit per photon limit.

---

[1] A. Kirmani, D. Venkatraman, A. Colaço, F. N. C. Wong, and V. K. Goyal, "High photon efficiency computational range imaging using spatio-temporal statistical regularization," in *Proc. CLEO*, San Jose, CA, Jun. 2013, paper QF1B.2.

[2] Z. T. Harmany, R. F. Marcia, and R. M. Willett, "Sparsity-regularized photon-limited imaging," in *IEEE Int. Symp. Biomed. Im.: From Nano to Macro*, Rotterdam, Apr. 2010, pp. 772–775.

[3] R. M. Gagliardi and S. Karp, *Optical Communications.* Wiley, 1976.

[4] G. Buller and A. Wallace, "Ranging and three-dimensional imaging using time-correlated single-photon counting and point-by-point acquisition," *IEEE J. Sel. Top. Quantum Electron.*, vol. 13, no. 4, pp. 1006–1015, Jul.–Aug. 2007.

[5] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory.* Upper Saddle River, NJ: Prentice Hall, 1993.

[6] S. B. Gokturk, H. Yalcin, and C. Bamji, "A time-of-flight depth sensor — system description, issues and solutions," in *Proc. Conf. Comput. Vis. Pattern Recog. Workshop*, 2004, p. 35.

[7] G. J. Iddan and G. Yahav, "3d imaging in the studio (and elsewhere)," in *Proc. SPIE*, vol. 4298, 2001, pp. 48–55.

[8] A. McCarthy, R. J. Collins, N. J. Krichel, V. Fernández, A. M. Wallace, and G. S. Buller, "Long-range time-of-flight scanning sensor based on high-speed time-correlated single-photon counting," *Appl. Optics*, vol. 48, no. 32, pp. 6241–6251, Nov. 2009.

[9] S. B. Kang, J. A. Webb, L. C. Zitnick, and T. Kanade, "A multibaseline stereo system with active illumination and real-time image acquisition," in *Proc. 5th Int. Conf. Comput. Vis.*, 1995, pp. 88–93.

[10] R. Garnett, T. Huegerich, C. Chui, and W. He, "A universal noise removal algorithm with an impulse detector," *IEEE Trans. Image Process.*, vol. 14, no. 11, pp. 1747–1754, Nov. 2005.

[11] A. Einstein, B. Podolsky, and N. Rosen, "Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?" "doibase ·10.1103/PhysRev.47.777 Phys. Rev. **47**, 777–780 (1935).

[12] J. C. Howell, R. S. Bennink, S. J. Bentley, and R. W. Boyd, "Realization of the Einstein-Podolsky-Rosen Paradox Using Momentum- and Position-Entangled Photons from Spontaneous Parametric Down Conversion," "doibase ·10.1103/PhysRevLett.92.210403 Phys. Rev. Lett. **92**, 210403 (2004).

[13] J. B. Pors, *Entangling Light in High Dimensions*, Ph.D. thesis, Leiden University (2011).

[14] A. K. Ekert, "Quantum Cryptography Based on Bell's Theorem," "doibase ·10.1103/PhysRevLett.67.661 Phys. Rev. Lett. **67**, 661–663 (1991).

[15] S. P. Walborn, D. S. Lemelle, M. P. Almeida, and P. H. Souto Ribeiro, "Quantum Key Distribution with Higher-Order Alphabets Using Spatially Encoded Qudits," "doibase·10.1103/PhysRevLett.96.090501 Phys. Rev. Lett. **96**, 090501 (2006).

[16] S. P. Walborn, D. S. Lemelle, D. S. Tasca, and P. H. Souto Ribeiro, "Schemes for Quantum Key Distribution with Higher-order Alphabets Using Single-photon Fractional Fourier Optics," "doibase·10.1103/PhysRevA.77.062323 Phys. Rev. A **77**, 062323 (2008).

[17] C. H. Bennett and S. J. Wiesner, "Communication via One- and Two-particle Operators on Einstein-Podolsky-Rosen States," "doibase·10.1103/PhysRevLett.69.2881 Phys. Rev. Lett. **69**, 2881–2884 (1992).

[18] S. L. Braunstein and H. J. Kimble, "Dense Coding for Continuous Variables," "doibase·10.1103/PhysRevA.61.042302 Phys. Rev. A **61**, 042302 (2000).

[19] T. B. Pittman, Y. H. Shih, D. V. Strekalov, and A. V. Sergienko, "Optical Imaging by Means of Two-photon Quantum Entanglement," "doibase ·10.1103/PhysRevA.52.R3429 Phys. Rev. A **52**, R3429–R3432 (1995).

[20] A. F. Abouraddy, B. E. A. Saleh, A. V. Sergienko, and M. C. Teich, "Role of Entanglement in Two-Photon Imaging," "doibase ·10.1103/PhysRevLett.87.123602 Phys. Rev. Lett. **87**, 123602 (2001).

[21] D. S. Tasca, R. M. Gomes, F. Toscano, P. H. Souto Ribeiro, and S. P. Walborn, "Continuous-variable Quantum Computation with Spatial Degrees of Freedom of Photons," "doibase ·10.1103/PhysRevA.83.052325 Phys. Rev. A **83**, 052325 (2011).

[22] S. P. Walborn, D. S. Ether, R. L. de Matos Filho, and N. Zagury, "Quantum Teleportation of the Angular Spectrum of a Single-photon Field," "doibase ·10.1103/PhysRevA.76.033801 Phys. Rev. A **76**, 033801 (2007).

[23] H. Di Lorenzo Pires, C. H. Monken, and M. P. van Exter, "Direct Measurement of Transverse-mode Entanglement in Two-photon States," ·http://dx.doi.org/10.1103/PhysRevA.80.022307 Phys. Rev. A **80**, 022307 (2009).

[24] V. C. Coffey, "Seeing in the Dark: Defense Applications of IR Imaging," ·http://dx.doi.org/10.1364/OPN.22.4.000026 Opt. Photonics News **22**, 26–31 (2011).

[25] M. A Albota, "Three-Dimensional Imaging Laser Radars with Geiger-Mode Avalange Photodiode Arrays," ·http://llwww.ll.mit.edu/publications/journal/pdf/vol13_no2/13_2laserradars3d.pdf Lincoln Laboratory Journal **13**, 351 − −367(2 10.1117/12.861600 Proc. SPIE, Infrared Remote Sensing and Instrumentation XVIII, Vol. 7808 (SPIE, 2010) p. 78080C.

[26] M. P. Edgar, D. S. Tasca, F. Izdebski, R. E. Warburton, J. Leach, M. Agnew, G. S. Buller, R. W. Boyd, and M. J. Padgett, "Imaging High-dimensional Spatial Entanglement with a Camera," ·http://dx.doi.org/10.1038/ncomms1988 Nat. Commun. **3** (2012), 10.1038/ncomms1988.

[28] P. B. Dixon, G A. Howland, J. Schneeloch, and J. C. Howell, "Quantum Mutual Information Capacity for High-Dimensional Entangled States," "doibase ·10.1103/PhysRevLett.108.143603 Phys. Rev. Lett. **108**, 143603 (2012).

[29] R. Bellman, *Dynamic Programming (Dover Books on Computer Science)* (Dover Publications, 2003).

[30] J. Schneeloch, P. B. Dixon, G. A. Howland, C. J. Broadbent, and J. C. Howell, ⁓http://arxiv.org/abs/1210.4234 "Witnessing Continuous Variable Entanglement with Discrete Measurements," (2012), arXiv, ⁓http://arxiv.org/abs/quant-ph/1210.4234 arXiv:quant-ph/1210.4234 .

[81] D. L. Donoho, "Compressed Sensing," "doibase-10.1109/TIT.2006.871582 Information Theory, IEEE Transactions on **52**, 1289 –1306 (2006).

[82] E. J. Candes and J. Romberg, "Sparsity and Incoherence in Compressive Sampling," ⁓http://stacks.iop.org/0266-5611/23/i=3/a=008 Inverse Problems **23**, 969 (2007).

[33] R.G. Baraniuk, "Compressive Sensing [Lecture Notes]," "doibase-10.1109/MSP.2007.4286571 Signal Processing Magazine, IEEE **24**, 118 –121 (2007).

[34] E. J. Candes and M. B. Wakin, "An Introduction To Compressive Sampling," "doibase-10.1109/MSP.2007.914731 Signal Processing Magazine, IEEE **25**, 21 –30 (2008).

[35] E. J. Candes and J. Romberg, http://users.ece.gatech.edu/˜justin/l1magic/ "l1-Magic: Recovery of Sparse Signals via Convex Programming," (2005).

[36] E.J. Candes, J. Romberg, and T. Tao, "Robust Uncertainty Principles: Exact Signal Reconstruction from Highly Incomplete Frequency Information," "doibase-10.1109/TIT.2005.862083 Information Theory, IEEE Transactions on **52**, 489 − 509 (2006).

[37] D. Gross, Y. K. Liu, S. T. Flammia, S. Becker, and J. Eisert, "Quantum State Tomography via Compressed Sensing," "doibase-10.1103/PhysRevLett.105.150401 Phys. Rev. Lett. **105**, 150401 (2010).

[38] A. Shabani, R. L. Kosut, M. Mohseni, H. Rabitz, M. A. Broome, M. P. Almeida, A. Fedrizzi, and A. G. White, "Efficient Measurement of Quantum Dynamics via Compressive Sensing," "doibase-10.1103/PhysRevLett.106.100401 Phys. Rev. Lett. **106**, 100401 (2011).

[39] O. Katz, Y. Bromberg, and Y. Silberberg, "Compressive Ghost Imaging," "doibase-10.1063/1.3238296 Applied Physics Letters **95**, 131110 (2009).

[40] P. Zerom, K. W. C. Chan, J. C. Howell, and R. W. Boyd, "Entangled-photon Compressive Ghost Imaging," "doibase-10.1103/PhysRevA.84.061804 Phys. Rev. A **84**, 061804 (2011).

[41] M.B. Wakin, J.N. Laska, M.F. Duarte, D. Baron, S. Sarvotham, D. Takhar, K.F. Kelly, and R.G. Baraniuk, "An Architecture for Compressive Imaging," in "doibase -·10.1109/ICIP.2006.312577 *Image Processing, 2006 IEEE International Conference on* (2006) pp. 1273 –1276.

[75] 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," "doibase-10.1109/MSP.2007.914730 Signal Processing Magazine, IEEE **25**, 83 –91 (2008).

[43] M.F. Duarte and R.G. Baraniuk, "Kronecker product matrices for compressive sensing," in "doibase -·10.1109/ICASSP.2010.5495900 *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on* (2010) pp. 3650 –3653.

[44] M.F. Duarte and R.G. Baraniuk, "Kronecker compressive sensing," "doibase-10.1109/TIP.2011.2165289 Image Processing, IEEE Transactions on **21**, 494 –504 (2012).

[45] C. Li, W. Yin, and Y. Zhang, ⁓http://www.caam.rice.edu/ optimization/L1/TVAL3/ "TVAL3: TV Minimization by Augmented Lagrangian and ALternating Direction ALgorithms," (2010).

[46] R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis* (Cambridge University Press, 1994).

[47] Rebecca M. Willett, Roummel F. Marcia, and Jonathan M. Nichols, "Compressed sensing for practical optical imaging systems: a tutorial," "doibase-10.1117/1.3596602 Optical Engineering **50**, 072601–072601–13 (2011).

[48] D.L. Donoho, A. Maleki, and A. Montanari, "The noise-sensitivity phase transition in compressed sensing," "doibase -·10.1109/TIT.2011.2165823 Information Theory, IEEE Transactions on **57**, 6920 –6941 (2011).

[49] Yihong Wu and S. Verdu, "Optimal phase transitions in compressed sensing," "doibase-10.1109/TIT.2012.2205894 Information Theory, IEEE Transactions on **58**, 6241 –6263 (2012).

[50] G. Reeves and M. Gastpar, "Compressed sensing phase transitions: Rigorous bounds versus replica predictions," in "doibase-10.1109/CISS.2012.6310927 *Information Sciences and Systems (CISS), 2012 46th Annual Conference on* (2012) pp. 1 –6.

[51] R.M. Willett and M. Raginsky, "Performance bounds on compressed sensing with poisson noise," in "doibase -·10.1109/ISIT.2009.5205258 *Information Theory, 2009. ISIT 2009. IEEE International Symposium on* (2009) pp. 174 –178.

[52] Z.T. Harmany, R.F. Marcia, and R.M. Willett, "Sparse poisson intensity reconstruction algorithms," in "doibase-10.1109/SSP.2009.5278495 *Statistical Signal Processing, 2009. SSP '09. IEEE/SP 15th Workshop on* (2009) pp. 634 –637.

[53] Surya Ganguli and Haim Sompolinsky, "Statistical mechanics of compressed sensing," "doibase -·10.1103/PhysRevLett.104.188701 Phys. Rev. Lett. **104**, 188701 (2010).

[87] M. A. T. Figueiredo, R. D. Nowak, and S. J. Wright, "Gradient Projection for Sparse Reconstruction: Application to Compressed Sensing and Other Inverse Problems," "doibase -·10.1109/JSTSP.2007.910281 Selected Topics in Signal Processing, IEEE Journal of **1**, 586–597 (2007).

[55] P.R. Gill, A. Wang, and A. Molnar, "The in-crowd algorithm for fast basis pursuit denoising," "doibase-10.1109/TSP.2011.2161292 Signal Processing, IEEE Transactions on **59**, 4595 –4605 (2011).

[56] S. M. Barnett and S. J. D. Phoenix, "Entropy as a Measure of Quantum Optical Correlation," "doibase -10.1103/PhysRevA.40.2404 Phys. Rev. A **40**, 2404–2409 (1989).

[57] S. P. Walborn, B. G. Taketani, A. Salles, F. Toscano, and R. L. de Matos Filho, "Entropic Entanglement Criteria for Continuous Variables," "doibase -10.1103/PhysRevLett.103.160505 Phys. Rev. Lett. **103**, 160505 (2009).

[58] S. P. Walborn, A. Salles, R. M. Gomes, F. Toscano, and P. H. Souto Ribeiro, "Revealing Hidden Einstein-Podolsky-Rosen Nonlocality," "doibase -10.1103/PhysRevLett.106.130402 Phys. Rev. Lett. **106**, 130402 (2011).

[59] M. V. Fedorov, Yu. M. Mikhailova, and P. A. Volkov, "Gaussian Modelling and Schmidt Modes of SPDC Biphoton States," -http://stacks.iop.org/0953-4075/42/i=17/a=175503 Journal of Physics B: Atomic, Molecular and Optical Physics **42**, 175503 (2009).

[60] M. D. Reid, "Demonstration of the Einstein-Podolsky-Rosen Paradox Using Nondegenerate Parametric Amplification," "doibase-10.1103/PhysRevA.40.913 Phys. Rev. A **40**, 913–923 (1989).

[61] Lu-Ming Duan, G. Giedke, J. I. Cirac, and P. Zoller, "Inseparability criterion for continuous variable systems," "doibase-10.1103/PhysRevLett.84.2722 Phys. Rev. Lett. **84**, 2722–2725 (2000).

[62] R. Simon, "Peres-horodecki separability criterion for continuous variable systems," Phys. Rev. Lett. **84**.

[63] I. Ali-Khan, C. J. Broadbent, and J. C. Howell, "Large-Alphabet Quantum Key Distribution Using Energy-Time Entangled Bipartite States," "doibase-10.1103/PhysRevLett.98.060503 Phys. Rev. Lett. **98**, 060503 (2007).

[64] W. Gong and S. Han, "Correlated Imaging in Scattering Media," "doibase-10.1364/OL.36.000394 Opt. Lett. **36**, 394–396 (2011)
.

[65] M.-C. Amann, T. Bosch, M. Lescure, R. Myllyla, and M. Rioux, "Laser ranging: a critical review of usual techniques for distance measurement," Optical Engineering **40**, 10–19 (2001).

[66] C. Mallet and F. Bretar, "Full-waveform topographic lidar: State-of-the-art," {ISPRS} Journal of Photogrammetry and Remote Sensing **64**, 1 – 16 (2009).

[67] S. Hussmann and T. Liepert, "Three-dimensional tof robot vision system," Instrumentation and Measurement, IEEE Transactions on **58**, 141–146 (2009).

[68] S. Foix, G. Alenya, and C. Torras, "Lock-in time-of-flight (tof) cameras: A survey," Sensors Journal, IEEE **11**, 1917–1926 (2011).

[69] B. Schwarz, "Mapping the world in 3d," Nature Photonics **4**, 429–430 (2010).

[70] A. McCarthy, N. J. Krichel, N. R. Gemmell, X. Ren, M. G. Tanner, S. N. Dorenbos, V. Zwiller, R. H. Hadfield, and G. S. Buller, "Kilometer-range, high resolution depth imaging via 1560 nm wavelength single-photon detection," Opt. Express **21**, 8904–8915 (2013).

[71] M. Richard and W. Davis, "Jigsaw: A foliage-penetrating 3d imaging laser radar system," Lincoln Laboratory Journal **15**, 1 (2005).

[72] M. Vaidyanathan, S. Blask, T. Higgins, W. Clifton, D. Davidsohn, R. Carson, V. Reynolds, J. Pfannenstiel, R. Cannata, R. Marino, J. Drover, R. Hatch, D. Schue, R. Freehart, G. Rowe, J. Mooney, C. Hart, B. Stanley, J. McLaughlin, E.-I. Lee, J. Berenholtz, B. Aull, J. Zayhowski, A. Vasile, P. Ramaswami, K. Ingersoll, T. Amoruso, I. Khan, W. Davis, and R. Heinrichs, "Jigsaw phase iii: a miniaturized airborne 3-d imaging laser radar with photon-counting sensitivity for foliage penetration," pp. 65500N–65500N–12 (2007).

[73] M. Entwistle, M. A. Itzler, J. Chen, M. Owens, K. Patel, X. Jiang, K. Slomkowski, and S. Rangwala, "Geiger-mode apd camera system for single-photon 3d ladar imaging," in "Proc. of SPIE Vol," , vol. 8375 (2012), vol. 8375, pp. 83750D–1.

[74] M. A. Itzler, M. Entwistle, M. Owens, K. Patel, X. Jiang, K. Slomkowski, S. Rangwala, P. F. Zalud, T. Senko, J. Tower *et al.*, "Comparison of 32 x 128 and 32 x 32 geiger-mode apd fpas for single photon 3d ladar imaging," in "SPIE Defense, Security, and Sensing," (International Society for Optics and Photonics, 2011), pp. 80330G–80330G.

[75] M. Duarte, M. Davenport, D. Takhar, J. Laska, T. Sun, K. Kelly, and R. Baraniuk, "Single-pixel imaging via compressive sampling," Signal Processing Magazine, IEEE **25**, 83–91 (2008).

[76] M. Sarkis and K. Diepold, "Depth map compression via compressed sensing," in "Image Processing (ICIP), 2009 16th IEEE International Conference on," (2009), pp. 737–740.

[77] A. Kirmani, A. Colaço, F. N. C. Wong, and V. K. Goyal, "Exploiting sparsity in time-of-flight range acquisition using a single time-resolved sensor," Opt. Express **19**, 21485–21507 (2011).

[78] G. A. Howland, P. B. Dixon, and J. C. Howell, "Photon-counting compressive sensing laser radar for 3d imaging," Appl. Opt. **50**, 5917–5920 (2011).

[79] L. Li, L. Wu, X. Wang, and E. Dang, "Gated viewing laser imaging with compressive sensing," Appl. Opt. **51**, 2706–2712 (2012).

[80] W. R. Babbitt, Z. W. Barber, and C. Renner, "Compressive laser ranging," Opt. Lett. **36**, 4794–4796 (2011).

[81] D. L. Donoho, "Compressed Sensing," Information Theory, IEEE Transactions on **52**, 1289 –1306 (2006).

[82] E. Cands and J. Romberg, "Sparsity and incoherence in compressive sampling," Inverse Problems **23**, 969 (2007).

[83] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse mri: The application of compressed sensing for rapid mr imaging," Magnetic Resonance in Medicine **58**, 1182–1195 (2007).

[84] J. Bobin, J.-L. Starck, and R. Ottensamer, "Compressed sensing in astronomy," Selected Topics in Signal Processing, IEEE Journal of **2**, 718–726 (2008).

[85] S. T. Flammia, D. Gross, Y.-K. Liu, and J. Eisert, "Quantum tomography via compressed sensing: error bounds, sample complexity and efficient estimators," New Journal of Physics **14**, 095022 (2012).

[86] G. A. Howland and J. C. Howell, "Efficient high-dimensional entanglement imaging with a compressive-sensing double-pixel camera," Phys. Rev. X **3**, 011013 (2013).

[87] M. A. T. Figueiredo, R. D. Nowak, and S. J. Wright, "Gradient Projection for Sparse Reconstruction: Application to Compressed Sensing and Other Inverse Problems," Selected Topics in Signal Processing, IEEE Journal of **1**, 586–597 (2007).

[88] D. Donoho and I. Johnstone, "Threshold selection for wavelet shrinkage of noisy data," in "Engineering in Medicine and Biology Society, 1994. Engineering Advances: New Opportunities for Biomedical Engineers. Proceedings of the 16th Annual International Conference of the IEEE," (1994), pp. A24–A25 vol.1.

[89] C. Li, W. Yin, and Y. Zhang, "Tval3: Tv minimization by augmented lagrangian and alternating direction algorithms," (2009).

[90] V. Cevher, A. Sankaranarayanan, M. F. Duarte, D. Reddy, R. G. Baraniuk, and R. Chellappa, "Compressive sensing for background subtraction," in "Computer Vision–ECCV 2008," (Springer, 2008), pp. 155–168.

[91] O. S. Magaña-Loaiza, G. A. Howland, M. Malik, J. C. Howell, and R. W. Boyd, "Compressive object tracking using entangled photons," Applied Physics Letters **102**, 231104 (2013).

[92] T. B. Pittman, Y. H. Shih, D. V. Strekalov, and A. V. Sergienko, "Optical imaging by means of two-photon quantum entanglement," Phys. Rev. A **52,** R3429–R3432 (1995).

[93] A. Valencia, G. Scarcelli, M. DAngelo, and Y. Shih, "Two-photon imaging with thermal light," Phys. Rev. Lett. **94,** 063601 (2005).

[94] F. Ferri, D. Magatti, A. Gatti, M. Bache, E. Brambilla, and L. A. Lugiato, "High-resolution ghost image and ghost diffraction experiments with thermal light," Phys. Rev. Lett. **94,** 183602 (2005).

[95] G. Scarcelli, V. Berardi, and Y. Shih, "Can two-photon correlation of chaotic light be considered as correlation of intensity fluctuations?," Phys. Rev. Lett. **96,** 063602 (2006).

[96] B. I. Erkmen and J. H. Shapiro, "Unified theory of ghost imaging with Gaussian-state light," Phys. Rev. A **77,** 043809 (2008).

[97] J. H. Shapiro, and R. W. Boyd, "The physics of ghost imaging. Quantum Inf. Process. **11,** 949–993 (2012).

[98] Y. Shih, "The physics of ghost imaging: nonlocal interference or local intensity fluctuation correlation?," Quantum Inf. Process. **11,** 995–1001 (2012).

[99] J. H. Shapiro, and R. W. Boyd, "Response to 'The physics of ghost imagingnonlocal interference or local intensity fluctuation correlation?,' Quantum Inf. Process. **11,** 1003–1011 (2012).

[100] N. D. Hardy and J. H. Shapiro, "Computational ghost imaging versus imaging laser radar for three-dimensional imaging," Phys. Rev. A **87,** 023820 (2013).

[101] J. H. Shapiro, "Computational ghost imaging," Phys. Rev. A **78,** 061802(R) (2008).

[102] O. Katz, Y. Bromberg, and Y. Silberberg, "Compressive ghost imaging," Appl. Phys. Lett. **95,** 131110 (2009).

[103] N. D. Hardy, "Improving optical remote imaging via graphical modeling," Ph.D. thesis proposal, MIT Dept. of Elect. Eng. and Comput. Sci., July 2013.

[104] J. H. Shapiro, J. Schneeloch, G. A. Howland, and J. C. Howell, "Laser radar point-target localization at high photon efficiency," Digest of CLEO 2013, (Opt. Soc. Am., Washington, DC, 2013) paper CTu1H.5

[105] D. Venkatraman, N. D. Hardy, F. N. C. Wong, and J. H. Shapiro, "Classical far-field phase-sensitive ghost imaging," Opt. Lett. **36,** 3684–3686 (2011).

[106] H. Ollivier and W. Zurek, "Quantum discord: a measure of quantumness of correlation," Phys. Rev. Lett **105,** 020503 (2010).

[107] S. Ragy and G. Adesso, "Nature of light correlations in ghost imaging," Sci Rep. **2,** 651 (2012).

[108] J. H. Shapiro, D. Venkatraman, and F. N. C. Wong, "Ghost imaging without discord," Sci. Rep. **3,** 1849 (2013).

[109] C. Zhao, W. Gong, M. Chen, E. Li, H. Wang, W. Xu, and S. Han, "Ghost imaging lidar via sparsity constraints," Appl Phys. Lett. **101,** 141123 (2012).

[110] S. Lauritzen, *Graphical models* (Oxford Univ. Press, Oxford, 1996).

[111] F. Jensen, *An introduction to Bayesian networks* (Springer-Verlag, Berlin, 1996).

[112] J. Pearl, *Probabilistic reasoning in intelligent systems: networks of plausible inference* (Morgan Kaufmann, San Francisco, 1988).

[113] F. Kschischang, B. J. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," IEEE Trans. Inform. Theory **47,** 498–519 (2001).

[114] D. L. Donoho, A. Maleki, and A. Montanari, "Message-passing algorithms for compressed sensing," Proc. Natl. Acad, Sci., USA **106,** 18914–18919 (2009).

[115] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," Digest 2001 IEEE International Sympos. on Inform. Theory 2168–2172 (IEEE, Piscataway, 2011).

[116] J. T. Parker, V. Cevher, and P. Schniter, "Compressive sensing under matrix uncertainties: An approximate message passing approach," Conference Record of 2011 Asilomar Conference on Signals, Systems and Computers, 804–808 (IEEE, Piscataway, 2011).

[117] J. H. Shapiro, B. A. Capron, and R. C. Harney, "Imaging and target detection with a heterodyne-reception optical radar," Appl. Opt. **20,** 3292–3313 (1981).

[118] C. W. Helstrom, *Quantum Detection and Estimation Theory* (Academic, New York, 1976).

[119] B. I. Erkmen and J. H. Shapiro, "Ghost imaging: from quantum to classical to computational," Adv. Opt. Photon. **2,** 405–450 (2010).

[120] D. N. Klyshko, Sov. Phys. JETP 76, 1131 (1988).

[121] A. V. Belinskii and D. N. Klyshko, Sov. Phys. JETP 78, 259 (1994).

[122] T. B. Pittman, Y. H. Shih, D. V. Strekalov, and A. V. Sergienko, Phys. Rev. A 52, R3429 (1995).

[123] R. S. Bennink, S. J. Bentley, R. W. Boyd, Phys. Rev. Lett. **89**, 113601 (2002).

[124] O. Katz, Y. Bromberg, and Y. Silberberg, Appl. Phys. Lett. **95**, 113110 (2009).

[125] V. Katkovnik, J. Astola, *J. Opt. Soc. Am. A*, **8**, 1556 (2012).

[126] B. Jack, J. Leach, J. Romero, S. Franke-Arnold, M. Ritsch-Marte, S. M. Barnett, and M. J. Padgett, Phys. Rev. Lett. **103**, 083602 (2009).

[127] E. J. Candès and M. B. Wakin, IEEE Signal Proc. Mag., 21-30 (March 2008).

[128] W. Gong, S. Han, *J. Opt. Soc. Am. A*, **29**, 1571 (2012).

[129] G. Molina-Terriza, L. Rebane, J. P. Torres, L. Torner, S. Carrasco, J. Eur. Opt. Soc. **2**, 07014 (2007).

[130] L. Torner, J. P. Torres, S. Carrasco, Opt. Exp. **13**, 873 (2005).

[131] D. S. Simon, A. V. Sergienko, *Phys. Rev. A*, **85**, 043825, 2012.

[132] L. Allen, M. Padgett, M. Babiker, Prog. Opt. 39, 291 (1999).

[133] J. P. Torres, A. Alexandrescu, L. Torner, Phys. Rev. A **68**, 050301(R) (2003).

[134] B. E. A. Saleh, A. F. Abouraddy, A. V. Sergienko, and M. C. Teich, Phys. Rev. A **62**, 043816 (2000).

[135] X. F. Ren, G. P. Guo, B. Yu, J. Li and G. C. Guo, J. Opt. B: Quantum Semiclass. Opt. **6**, 243 (2004).

[136] R. S. Bennink, S. J. Bentley, R. W. Boyd, J. C. Howell, Phys. Rev. Lett. **92**, 033601 (2004).

[137] A. Gatti, E. Brambilla, M. Bache, L. A. Lugiato, Phys. Rev. Lett. **93**, 093602 (2004).

[138] A. Valencia, G. Scarcelli, M. D'Angelo, Y.H. Shih, Phys. Rev. Lett. **94**, 063601 (2005).

[139] G. Scarcelli, V. Berardi, Y. H. Shih, Phys. Rev. Lett. **96**, 063602 (2006).

[140] F. Ferri, D. Magatti, A. Gatti, M. Bache, E. Brambilla, L. A. Lugiato, Phys. Rev. Lett. **94**, 183602 (2005)

[141] E. J. Candès, Proc. Int. Cong. Mathematicians, Madrid, Spain (2006).

[142] E. Candès, J. Romberg, and T. Tao, IEEE Trans. Inform. Theory **52**, 489 (2006).

[143] E. Candès and T. Tao, IEEE Trans. Inform. Theory **52**, 5406 (2006).

[144] D. Donoho, IEEE Trans. Inform. Theory **52**, 1289 (2006).

[145] N. Uribe-Patarroyo, A. Fraine, D. S. Simon, O. Minaeva, A. V. Sergienko, *Phys. Rev. Lett.* **110**, 043601 (2013).

[146] A. Mair, A. Vaziri, G. Weihs, A. Zeilinger, Nature **412**, 313(2001).

[147] H. Di Lorenzo Pires, H. C. B. Florijn, M. P. van Exter, *Phys. Rev. Lett.*, **104**, 020505 (2010).

[148] J. Leach, B. Jack, Barry, J. Romero, A. K. Jha, K. Anand, A. M. Yao, S. Franke-Arnold, D. G. Ireland, R. W. Boyd, S. M. Barnett, M. J. Padgett, *Science*, **329**, 662 (2010).

[149] C. A. Fitzpatrick, D. S. Simon, A. V. Sergienko, in preparation (2013).

[150] A. M. Yao and M. J. Padgett, Adv. in Opt. & Phot. **3**, 161(2011).

[151] S. Franke-Arnold, L. Allen, M. Padgett, Laser & Photon. Rev. **2**, 299 (2008).

[152] L. Allen, M. W. Beijersbergen, R. J. C. Spreeuw, and J. P. Woerdman, Phys. Rev. A 45, 8185 (1992).

[153] D.V. Strekalov, A. V. Sergienko, D. N. Klyshko, and Y. H. Shih, Phys. Rev. Lett. 74, 3600 (1995).

[154] B Jack, P Aursand, S Franke-Arnold, D G Ireland, J Leach, S M Barnett, and M J Padgett, J. Opt. **13**, 064017 (2011).

[155] M. J. Padgett, L. and Allen, *Opt. Comm.*, **121** 36 (1995).