Quantum-inspired computational imaging

The rise of quantum-inspired computational imaging

Authors: Yoann Altmann[1], Stephen McLaughlin[1], Miles Padgett[2], Vivek K. Goyal[3], Alfred Hero[4], Daniele Faccio[2]*,

Affiliations:

[1]School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, UK.

[2]School of Physics & Astronomy, University of Glasgow, Glasgow, UK.

[3]Dept. of Electrical and Computer Engineering, Boston University, Boston, USA.

[4]Dept. of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, USA.

*Correspondence to: Daniele.Faccio@glasgow.ac.uk

Abstract:

Computational imaging combines measurement and computational methods with the aim of forming images even when the measurements are weak, few in number, or highly indirect. Although this field of research is nearly as old as photography itself, the recent surge in new quantum-inspired imaging sensors, together with a new wave of algorithms allowing on-chip, scalable and robust data processing, have induced a remarkable increase of activity with stunning results in the domain of low-light flux imaging and sensing. In this review, we provide an overview of the major challenges encountered in low-illumination (e.g., ultrafast) imaging and how these problems have recently been addressed for imaging applications in extreme conditions. These methods provide clear examples of the future imaging solutions to be developed, where the best results are expected to arise from an efficient co-design of the sensors and data analysis tools.

One Sentence Summary:

Cutting-edge single-photon detectors and advanced computational methods combine to create new imaging solutions and push back the boundary of the invisible.

Introduction: Computational imaging is the fusion of computational methods and imaging techniques with the aim of producing better images, where "better" has a multiplicity of meaning. A recent surge in new imaging sensors and in particular with single-photon sensors, combined with a new wave of computational algorithms, data handling capability and deep learning, has resulted in a remarkable surge of activity in the field.

One clear trend is a shift away from increasing the number of mega-pixels towards fusing camera data with computational processing and, if anything, decreasing the number of pixels, potentially to a single pixel. The incoming data may therefore not actually look like an image in the conventional sense but is transformed into one after a series of computational steps and/or modelling of how the light travels through the scene or the camera. This additional layer of computational processing frees us from the chains of conventional imaging techniques and removes many limitations in our imaging capability.

We briefly describe some of the most recent developments in the field including full 3D imaging of scenes that are hidden (e.g. around a corner or behind an obstacle), high resolution imaging with a single-pixel detector at wavelengths for which no cameras exist, cameras that can see through fog or inside the human body or cameras that mimic the human eye by creating detail only in areas of interest. We will also discuss how multispectral imaging using single-photon detectors can improve 3D reconstruction and provide richer information about the scene.

In a more general context, we discuss how single-photon detection technologies are transforming imaging capabilities with single-photon sensitive cameras that can take pictures at the lowest light levels and with the ability to create videos reaching a trillion frames per second. This improvement has enabled the capture of images of light beams travelling across a scene and provided opportunities to observe image distortions and peculiar relativistic temporal inversion effects that are due to the finite speed of light. The ability to capture light in flight also underpins some of the applications mentioned above, for example the ability to view a 3D scene from around a corner. Probabilistic modelling of the particle-like nature of light when using single-photon detectors has stimulated the birth of new computational techniques such as "first-photon imaging," which hints at the ultimate limits of information to be gained from detecting just one photon.

Single pixel and ghost imaging: Although most imaging techniques that have emerged recently are based on classical detectors and cameras, some of these approaches have been inspired by or have a tight connection with similar ideas in quantum imaging. A prime example is ghost imaging *(1)*, originally thought to be based purely on quantum principles but now recognized as being dependent upon spatial correlations that can arise from both quantum and classical light *(2)*. The realization that this technique does not require quantum light led to a merging of the fields of computational ghost imaging *(3)* with work on single-pixel cameras *(4)* and to an overall increase of activity in the field. In its quantum version, ghost imaging refers to the use of parametric down-conversion to create pairs of photons with correlated positions. If we detect the position of one photon with a standard camera and illuminate an object with the other position-correlated photon, it is sufficient to detect only the reflectance or transmittance of the object with a single-pixel detector, i.e. measure the correlation count between the beams to then reconstruct a full image by repeating the measurement with many different photon pairs (each of which will be randomly distributed due to the random nature of the correlated photon pair generation process) *(5-6)*. It is now acknowledged that the quantum properties of the correlated photons play no role in the image reconstruction process: thermal light split into two beams using a beam splitter can be used equally effectively, albeit at a higher photon flux *(7)*. Rather than using a beam splitter, it is possible to use a spatial light modulator to create a pattern where the copy is simply the computer memory. This approach therefore no longer requires a camera of any kind in the setup: the computer-generated pattern is already known and the image, I, can be retrieved by multiplying the single pixel readout, a_i, with the corresponding pattern, H_i, and then summing over all i patterns, i.e. $I = \sum_i a_i H_i$. This opens the route to so-called compressed single-pixel imaging: where assumptions about the spatial correlations of the image enable patterns to be judiciously chosen to require a far fewer number of patterns than the final number of image pixels, with compression factors up to 80-90%. This concept is not dissimilar from standard jpeg compression, which assumes that typical images are concentrated in their spatial frequencies, with the difference that now the compression is applied at the image acquisition stage. By this compression, single-pixel imaging is therefore transformed from a slow, relatively inefficient process into a highly efficient imaging technique that can operate at video frame rates in full colour *(8)*. More recent developments include the

extension to full 3D images where depth information is obtained by also using time-of-flight information *(9-10)*, i.e. alongside the object reflectivity one also keeps track of distance, d, from the temporal shift, τ, of the detected signal, as the two are related by the speed of light, $d = c\tau$. In general, this single-pixel technique suffers from having lower resolution and providing poorer quality images even when compared to a cell-phone camera. This limitation that can been partly overcome by taking inspiration from nature and implementing computational algorithms so that the system increases the density of the projected spatial patterns only in areas of interest, therefore increasing the spatial resolution in regions where it is needed and leaving the surrounding areas relatively less defined *(11)*. This is just one example of where computational techniques have fused together with the detection technology to provide more efficient sensing solutions. Another example is the "first-photon imaging" approach that emerged from a unique co-design of hardware and computational algorithms, built around the concept of single-photon detection.

First-photon imaging: An important legacy of recent interest in the field of quantum information science is the development of a series of detector technologies for single photons. The workhorse for most laboratories is the single-photon avalanche diode (SPAD). SPADs are in essence, semiconductor diodes that are reverse-biased beyond their breakdown threshold: a single photon (or even a thermally generated charge in the diode) is sufficient to lead to the rapid charge multiplication process (or avalanche) that creates a spike in the output current. A quenching mechanism stops the avalanche process before the diode is irreversibly damaged, leading also to a dead time during which the diode is insensitive to incident photons before being re-activated. The particle-like nature of a photon is revealed through the very short burst in the SPAD output current that can then be very precisely timed when a reference signal is also provided. The ability to precisely detect the photon arrival time can be used for long distance, high-precision LIDAR ranging: a distant object is illuminated with a point-like pulsed laser beam. Each outgoing pulse starts a counter, which is then stopped at time τ when a return photon is detected; accounting for the two directions of the light travel, the distance of the illuminated object is simply $d = c\tau/2$. Scanning the scene using this time-correlated single photon counting (TCSPC) technique can therefore provide a full 3D image (or depth image) of the scene *(12-14)*. However, TCSPC-based imaging can require very long acquisition times, in particular when photons return to the detector at low rate. Conventional processing techniques require: (i) operation in the photon-starved regime, i.e. 10% or less of the outgoing laser pulses should give rise to a detected return photon so that bias from detector dead times is negligible; and (ii) measurement over many illumination repetition periods so that 100-1000 photons or more are detected for each position. Under these conditions a faithful measurement of the photon arrival time is obtained. This approach can easily lead to acquisition times of a complex scene that can be of the order of many seconds or even minutes.

The computational imaging philosophy enables a significant reduction in the number of detected photons needed for 3D imaging *(15)*. In the "first-photon imaging" approach, only the very first detected photon at each scan location is used, so the acquisition time is limited primarily by the speed of scanning, and any detector dead time coincides with the scanning *(16)*. As shown in Fig. 2, using the number of pulses until a photon is detected as an indirect measurement of reflectivity along with a piecewise-smooth assumption for both reflectivity and depth, after several computational steps, a 3D image of a scene is produced as shown in Fig. 2. This approach builds a strong link between the computational steps and the detailed mechanism of single-photon detection, with various aspects such as the noise background and the particle-like nature of the

photons and their detection built into the information used to retrieve high-quality 3D images. Similar extreme photon efficiency can be achieved with a fixed dwell time at each scene position *(17)*, and principled statistical techniques for adapting the local spatial resolution to characteristics of the data enables background noise 25 times stronger than the back-reflected signal to be mitigated *(18)*. Additional performance improvements have been obtained with deep learning. Using an array of SPADs parallelizes the data acquisition and thus can increase imaging speed, though an array has coarser time resolution, translating to coarsened longitudinal distance measurements *(19)*. Methods for arrays can also be highly photon efficient *(20)*.

Non-Line of sight imaging: Photon counting has strongly impacted the field of Non-Line-Of-Sight (NLOS) imaging, i.e. the imaging objects that are for example hidden behind a wall, corner or obstacle *(21-30)*. Access to very high temporal resolution imaging systems allows reconstruction of a full 3D image of the hidden scene, as conceptually explained in Fig. 3(A). A short laser pulse is sent to a scattering surface chosen so as to scatter light behind the obstacle and thus illuminate the hidden scene. The hidden scene will reflect back a return echo that will once again hit the first scattering wall and return to the imaging system. An intuitive understanding of the hidden object reconstruction is based on the fact that the locus of points that can give rise to a back-scattered signal from a laser spot at a position $r_l = (x_l, y_l)$ and measured at a given point $r_i = (x_i, y_i)$ on the wall is given by $|r - r_l| + |r - r_i| + |z| = c\tau$. This equation describes an ellipsoid of points that can be recalculated for each detection point on the wall: each of these ellipsoids will overlap only at the points of origin of the (hidden object) scattering. Therefore, by summing over all ellipsoids, one obtains a high "intensity" (proportional to the overlap) in correspondence to the hidden object. With sufficient temporal resolution and additional processing to sharpen the retrieval, it is possible to reconstruct full 3D shapes: for example, 100 ps is sufficient to resolve cm-sized features. Practically, most retrieval techniques aim at iteratively finding a solution for $I = A\rho$ where $\rho(x, y, z)$ represents the physical distribution of the hidden object and the measured transient image intensity is

$$I(x, y, \tau) = \iiint \frac{1}{r_l^2 r_i^2} \delta(|r - r_l| + |r - r_i| + |z| - c\tau)\rho(x, y, z)\, dxdydz$$

The first demonstration of this technique was obtained using a streak camera *(21)* that provides very high 1-10 ps temporal resolution but at the expense of relatively long acquisition times, see Fig.3(B) and (C). Successive demonstrations resorted to single photon counting to reconstruct 3D images *(22)* and to perform tasks such as tracking of moving objects *(25)* and humans even over very large distances (more than 50 meters between the scattering wall and the imaging system *(26)*) using also single pixel SPADs. Recent improvements have demonstrated acquisition times of the order seconds for a full 3D scene reconstruction by modifying the acquisition scheme (photons are now collected coaxially, i.e. along the same (exact) trajectory as the outgoing the laser beam *(29)*) as a result of which measurement integral is simplified to

$$I(x, y, \tau) = \iiint \frac{1}{r^4} \delta(r + r_l - \tau c)\rho(x, y, z)\, dxdydz$$

where the radiometric term $1/r^4$, is now only a function of τ and can thus be removed from the integral. Overall, the result of this is that $I(x, y, \tau)$ reduces to a convolution that significantly reduces the computational retrieval times, paving the way to real-time reconstruction of 3D scenes.

This is an example of how imaging hardware and computational techniques have co-evolved to create a new imaging capability. It is worth pointing out that recent measurements have shown not only real-time capability, but also long-distance and full daylight operation *(26, 29)*, thus moving from proof-of-principle to first steps towards deployment in real-world applications in just a few years. An interesting challenge for this field of research starts from the observation that much of the technology involved in NLOS imaging is common with standard, direct line-of-sight of LIDAR, i.e. 3D imaging of environments using laser pulse time-of-flight measurements. In this sense, NLOS imaging has the potential to become a natural extension of LIDAR. In this context, there are clear applications for NLOS imaging, when combined with LIDAR, for urban safety and un-manned vehicles. It is also worth noting that future NASA missions will employ SPAD arrays for LIDAR mapping of planet surfaces and studies are currently underway to evaluate the potential of NLOS imaging to remotely (e.g. from a satellite) assess the internal structure of underground caves on planets in view of future human colonisation activities *(31)*.

Enhanced SPAD arrays for imaging in scattering media: Over the past several years a number of industrial and academic research groups have been developing a new generation of cameras where each individual pixel consists of a SPAD and all the TCSPC electronics and can thus efficiently acquire transient images with a resolution as low as 50 ps *(19, 32-34)*. This effectively implies that when combined with a high repetition rate laser for the active illumination of the scene, videos rates reaching up to 20 billion frames per second can be achieved *(35)*. This remarkable performance is probably better appreciated when expressed in terms of the actual imaging capability. At such frame rates per second, light propagates just 1.5 cm between successive frames implying that it is possible to actually freeze light in motion in much the same way that standard high-speed cameras can freeze the motion of a supersonic bullet. The first images of "light in flight" were shown in the late 1960's using nonlinear optical gating methods *(36-38)* but the first camera-based measurements were only demonstrated several years ago using a streak camera *(39)*. More recent measurements based on SPAD arrays have allowed the first capture of light pulses propagating in free space with total acquisitions times of the order of seconds or less *(35)*. SPAD array cameras have also been used to directly image laser pulse propagation through optical fibres: beyond the direct applications of these measurements e.g. to stand-off detection of the fibre parameters, these measurements combined a fusion of single photon data together with hyperspectral imaging over several different wavelengths (discussed also below) and computational processing through which the 32x32 pixel resolution was successfully up-sampled by using the temporal axis to re-interpolate the pulse trajectory in the (x,y) spatial plane *(40)*.

The ability to capture simultaneously spatial and high-resolution temporal information at very low light levels with SPAD cameras has recently been applied to other challenging imaging problems, for example to imaging and sensing through scattering and turbid media. Pavia *et al.* have for example, applied inverse retrieval methods in combination with spatial and temporal information from a linear SPAD array for tomographic reconstruction of objects hidden in murky water *(41)*. More recently, Heshmat et al. have acquired data with a planar SPAD array and reconstructed various shapes of objects obscured by a thick tissue phantom *(42)*. Their technique was called "All Photons Imaging", directly underlining the importance of the photon time-of-flight information that is recorded by the single-photon camera. We note that such approaches do not necessarily explicitly distinguish between the physical origin of the data sets. For example, temporal or spatial information are placed on equal footing and enter in the retrieval process in a manner where it is simply the quality of data that constrains the problem. This approach hints towards future

developments where the retrieval problem is further constrained by adding measurements from as many different sources and detectors as possible. In the broad regime of strong scattering, the camera will typically record an indistinct, diffuse illumination transmitted through the medium or reflected from the scene with little or no obvious information about any objects hidden behind or inside the scattering medium. Computational techniques are thus required in order to actually retrieve details about the object. This field of research is of particular interest for a number of applications such as medical diagnostics and imaging and sensing/imaging through fog.

With the rise autonomous underwater vehicles (AUVs), unmanned aerial vehicles (UAVs), robots or cars, water and fog are particularly challenging media for imaging systems. While sonar is a well-established technology for long range underwater imaging, it can suffer from low spatial resolution limited by the physics of sound propagation in the medium. While high power optical solutions can be used for short range imaging in relatively clear water, the presence of underwater scatterers between the active imaging system and the scene (e.g., the seabed) usually produce large quantities of reflected photons that can mask the returns from the scene of interest. Using a pulsed illumination source combined with sensitive single photon detectors, it is possible to discriminate the photons reflected due to scattering in the water from those (an extremely small refraction) that actually reach the surfaces of interest. For instance, in *(43)*, the authors demonstrated the ability to image underwater up to 8 attenuation lengths. When combining this edge-cutting technology with advanced statistical methods inspired by *(44)*, significant performance improvements could be achieved in terms of 3D reconstruction and estimation of the surface reflectivity by accounting for the distance-induced signal loss *(45)*. Reconstruction of 3D of terrestrial scenes at long distances suffers from similar limitations. Even if the measurements are performed under favorable conditions (e.g., dry conditions), the recorded signals can be significantly affected by atmospheric turbulence *(46-48)*, and solar illumination *(49, 50)*. Again, significant improvements in terms of detection accuracy *(51)* and maximal observable range *(52)* have been obtained by using adapted computational tools. The problem becomes even more acute in the presence of fog, which is a major concern for the next generation of automated cars. It has been demonstrated that it is technically possible to detect and analyze fog patches over long distances, provided that the laser power is sufficient to ensure a non-zero probability of photon reflection and a long enough acquisition time *(53, 54)*. In the automotive context, where the acquisition time is intrinsically limited by the vehicle displacement, more robust and computationally efficient strategies have been recently proposed *(55, 56)* and it is clear that the future solutions will embed physical modelling of both the medium and the data collection process within the information extraction.

Multispectral single-photon imaging. Multispectral and hyperspectral imaging, which are extensions of classical color (RGB) imaging, consist of imaging a scene using multiple wavelengths (from four to several hundreds or even thousands in hyperspectral images). These modalities have benefited from a significant body of research over more than 35 years, from the data collection community (57-59) and, more importantly, from the data processing and analysis community (60-65). Indeed, such modalities can be associated with a wide variety of computational problems, ranging from image acquisition (compressive sampling), restoration (denoising/deblurring, super-resolution), segmentation (classification) to source separation (spectral unmixing), object/anomaly detection and data fusion (e.g., so called pansharpening). While the main applications using (passive) multispectral imaging are in Earth and space observation, the proven benefits of imaging with multiple wavelengths simultaneously have

enabled its application in the food industry *(58, 66)* and a broader range of applications such as diagnostic dermatology *(67, 68)*. Active multispectral imaging is less sensitive to ambient illumination than passive imaging, which requires data acquisition under daylight condition (e.g., for Earth observation). Without fast timing capabilities, multi/hyperspectral imagers are however only able to provide 2D intensity profiles and are thus not adapted to analyze multi-layered 3D structures, such as forest canopies. Multispectral Lidar is a promising modality allowing for joint extraction of geometric (as single wavelength Lidar) and spectral (as passive multispectral images) information from the scene, while avoiding data registration issues potentially induced by the fusion of heterogeneous sensors. In *(69)*, it was demonstrated that it was possible to use multispectral single photon Lidar (MSPL) to remotely infer the spatial composition (leaves and branches) and the health of trees using only 4 wavelengths. More recently, new experiments have been designed to image up to 33 wavelengths (500-820nm) in free space *(70)* and 16 wavelengths underwater *(71)*. As a consequence, we have witnessed the development of algorithms inspired from passive hyperspectral imagery (3D datacubes) for analysis of MSPL data (4D datacubes). For instance, Bayesian methods have been proposed to cluster, in an unsupervised manner, spectrally similar objects while estimating their range, from photon-starved MSPL data *(72)*. This work was further developed in *(73, 74)* to classify pixels based on their spectral profiles in photon-starved regimes down to 1 photon per pixel and per spectral band on average (see Fig. 4). Such results are only possible by efficiently combining a highly sensitive raster scanning single-photon system allowing for submillimetre range resolution with hierarchical Bayesian models able to capture the intrinsic, yet faint, structures (e.g., spatial and spectral correlations) of the data. A significant improvement has been demonstrated using simulation methods (see next section) to reconstruct scenes (range and reflectivity profiles) with as few as 4 photons per pixel (with 4 spectral bands and one photon per pixel on average) *(73)*. Spectral unmixing presents another challenging problem encountered when analyzing multi/hyperspectral data, that is, when it is necessary to identify and quantify the materials/components present in the observed scene. Spectral unmixing thus generalizes classification methods by accounting for the fact that several "mixed" materials can be observed in a given pixel. Spectral unmixing methods allow for sub-pixel material quantification, which is particularly interesting for long-range imaging scenarios where the divergence of the laser beam cannot be neglected. In *(70)*, a novel computational method is developed for quantifying and locating 15 known materials from MSPL data consisting of 33 spectral bands, while detecting additional (potentially unknown) materials present in the scene. Again, this work demonstrated the possibility of material quantification and anomaly detection with as few as 1 photon per pixel and per spectral band, on average. It also illustrated how Bayesian modelling can be used for uncertainty quantification, e.g., for providing confidence intervals associated with estimated range profiles. As mentioned above, while the most recent single-photon detectors are very attractive for their high temporal resolution, their use to extract information from large scenes is currently limited by long acquisition times associated with raster scanning strategies. This is particularly limiting when several wavelengths are acquired in a sequential manner. To address this problem, compressive sampling strategies have been investigated to achieve faster MSLP data acquisition *(75, 76)*. Although computational methods adapted for scanning systems have been proposed, whereby a random number of spectral bands can be probed in a given pixel, the most promising results have been obtained with a simulated mosaic filter (4 wavelengths) whose actual implementation within a SPAD array in the near future should allow for the simultaneous acquisition of multiple pixels and fast reconstruction of range and reflectivity profiles. These results show how advanced computational methods can be used to enhance

information extraction from imaging systems but also improve the design of future detectors and detector arrays.

Computational methods in the photon starved regime. Generally speaking, from a computational perspective, extracting information from a scene can be seen as inferring a set of parameters $\boldsymbol{x}$, which may have a physical meaning (or not), for a set of measurements $\boldsymbol{y}$ recorded by an imaging/sensing device. These parameters can take continuous values, (e.g., light field intensities, object positions and velocities) or discrete values, (e.g., the number objects, binary values representatives of the presence or absence of objects). Two main families of methods can be adopted to design algorithms for data analysis, namely, supervised machine learning approaches and statistical signal processing approaches, although hybrid methods can also be used. The choice of the most suitable approaches depends primarily on the complexity of the computational model, as well as the computational budget available, i.e., the expected processing time, data storage limitations and the desired quality of the information extracted. Supervised machine learning (including deep learning) approaches are particularly well suited for applications where a sufficient quantity of ground truth data/information or reference data is available *(77-80)*. Such methods rely on a two-stage process, referred to as the training stage and the test stage. Starting from a forward model $\boldsymbol{y} \approx g(\boldsymbol{x})$, relating the measured data $\boldsymbol{y}$ to the unknown source parameters $\boldsymbol{x}$, the training stage uses a set of measurements and corresponding parameters to learn the inverse mapping $h(\cdot)$ between the measurements and the set of parameters to be recovered, i.e., it fits an inverse model $\boldsymbol{x} \approx h(\boldsymbol{y})$. In contrast to model-based statistical methods, data-driven machine learning approaches do not rely on the knowledge a forward model $g(\cdot)$. Thus, these methods can often be applied to complex problems where the forward model is unknown or too complicated to be derived analytically but there is plenty of training data available. The training stage controls the quality of the estimation of the mapping $h(\cdot)$, which in turn depends on the representational power of the machine learning algorithm, and on the quality/diversity of the training samples. Machine learning approaches have been successfully applied to imaging applications such as imaging through multimode fibres *(77)*, lensless imaging of phase objects *(78)* and identification of human pose from behind behind a diffusive screen *(79)*. SPAD cameras have been specifically applied to the problem if identifying both the position and identity of a people hidden behind a wall *(80)* scattering walls *(81)* or camouflage *(82)*. Their use can become limited due to their poor adaptability properties (e.g., they have to be re-trained for different acquisition scenarios) and the lack of ground truth information as is often the case in photon-starved regimes where the goal is to extract information previously inaccessible.

In such scenarios, model-based statistical methods can be more attractive, as they can better fit the data to a forward model $g(\cdot)$. Physical considerations, such as light transport theory through the medium and the detector can often guide the choice of the most appropriate models, although non-physically inspired approximations can be used to make the model fitting algorithm more computationally tractable. When there is measurement uncertainty and noise the forward model can be better characterized by the conditional probability distribution $f(\boldsymbol{y}|\boldsymbol{x})$, which describes the statistical variation of the measurements $\boldsymbol{y}$ for a given source parameter value $\boldsymbol{x}$. For fixed value $\boldsymbol{y}$, the function $l_y(\boldsymbol{x}) = f(\boldsymbol{y}|\boldsymbol{x})$, called the likelihood function, quantifies the likelihood that the source value $\boldsymbol{x}$ generated the observed value $\boldsymbol{y}$. The maximum likelihood principle forms an estimate of $\boldsymbol{x}$ from $\boldsymbol{y}$ by maximizing the likelihood over $\boldsymbol{x}$. However, the maximum likelihood estimate is often not unique in high dimensional inverse problems such as imaging. However, often additional information about $\boldsymbol{x}$ is available, e.g., a priori knowledge about positivity, smoothness,

or sparsity, and can be used to improve on the maximum likelihood estimate. Let ϕ be a regularization function such that $\phi(\boldsymbol{x})$ is small when $\boldsymbol{x}$ complies with a priori knowledge and is large otherwise. Then it is possible to recover $\boldsymbol{x}$ by minimizing the cost function $C_{\boldsymbol{y}}(\boldsymbol{x}) = -\log\left(l_y(\boldsymbol{x})\right) + \phi(\boldsymbol{x})$, i.e.,

$$\min_{\boldsymbol{x}} C_{\boldsymbol{y}}(\boldsymbol{x}) = \min_{\boldsymbol{x}} \left\{-\log\left(l_y(\boldsymbol{x})\right) + \phi(\boldsymbol{x})\right\}.$$

If $\phi(\boldsymbol{x})$ can be associated with a proper density $f(\boldsymbol{x}) \propto e^{-\phi(\boldsymbol{x})}$, called the prior distribution, this penalized likelihood estimation strategy can be interpreted in the Bayesian formalism as a maximum a posteriori estimation procedure, i.e., the above minimization is equivalent to maximizing the posterior density of $\boldsymbol{x}$

$$f(\boldsymbol{x}|\boldsymbol{y}) = f(\boldsymbol{y}|\boldsymbol{x})f(x)/f(y),$$

where $f(\boldsymbol{y})$ is a density that does not depend on $\boldsymbol{x}$. These and related likelihood based approaches have been adopted by many researchers in low photon imaging, e.g., (16, 17, 75).

The Bayesian formalism allows for the observed data to be combined with additional information available in a principled manner. This also allows so-called a posteriori measures of uncertainty to be derived. However, such measures cannot be computed analytically in most practical applications because of the complexity of good spatial correlation models, so they often rely on high-dimensional integrations. While significant advantage may be gained from computationally simple pixel-by-pixel adaptation *(83)*, a classical approach thus consists of approximating these measures (e.g., a posteriori variances/covariances or confidence intervals) using variational approximations or simulation techniques. Markov Chain Monte Carlo (MCMC) methods are particularly well adapted for inference in difficult scenarios where the cost function or posterior distribution of interest has multiple modes and multiple solutions potentially admissible. For instance, such methods have been successfully applied to object detection *(51)*, material quantification and anomaly detection *(70)* from low-flux single-photon Lidar measurements.

Conclusions. With the rapid advance in imaging cameras and sensors together with a leap forward in computational capacity, we see enormous potential for significant innovation over the next several years. The main challenge—or the main opportunity—at this stage is the co-development of sensors and computational algorithms that are built around the physical processes of the photon transport and detection processes. We have provided examples showing progress in this direction, ranging from first-photon imaging techniques to photon-starved hyperspectral imaging. The trend seen in commercial cameras between 2000-2015, characterized by a constant drive towards higher pixel counts, has slowly subsided, giving way to a different approach whereby both performance and functionality are increased by combining multiple sensors through computational processing. Obvious examples are recent advances in cell-phone technology, arguably one of the main drivers behind imaging technology, that now boasts multiple cameras and lenses providing depth perception, improved signal to noise ratio and other functionalities such as 3D face recognition. With SPAD cameras also gradually making their appearance on the commercial scene, single-photon imaging and computational techniques offer an exciting avenue for future innovation in situations where previously imaging was not thought to be possible. We have briefly discussed

examples such as imaging though denser scattering media such as the human body or fog, and full 3D imaging of scenes that are around a corner or beyond the direct line of sight, e.g., around a corner. These examples are particularly relevant in demonstrating the progress that can be made when the photon transport models and computational approaches are integrated with the new generation of photon detectors. The expectation is that over the next several years we will witness a remarkable growth of computational imaging methods, driven by and also driving new technologies such as single-photon SPAD arrays that will revolutionize nearly every aspect of human activity, ranging from medical diagnostics to urban safety and space missions.

References

(1) T. B. Pittman, Y. H. Shih, D. V. Strekalov, A. V. Sergienko, Optical imaging by means of two-photon quantum entanglement. *Phys Rev A* **52**(2), 3429–3432 (1995).

(2) J. H. Shapiro, R. W. Boyd, The physics of ghost imaging. *Quantum Inf. Process.* **11**(4), 949–993 (2012) [doi:10.1007/s11128-011-0356-5].

(3) J. Shapiro, Computational ghost imaging. *Phys Rev A* **78**(6), 061802 (2008) [doi:10.1103/PhysRevA.78.061802].

(4) M. F. Duart, M. A. Davenport, T. Dharmpal, J. N. Laska, S. Ting, K. F. Kelly, R. G. Baraniuk, Single-pixel imaging via compressive sampling. *IEEE Signal Process. Mag.* **25**(2), 83–91 (2008) [doi:10.1109/MSP.2007.914730].

(5) D. Shin, J. H. Shapiro, V. K. Goyal, Performance analysis of low-flux least-squares single-pixel imaging. *IEEE Signal Process. Lett.* **23**(12), 1756-1760 (2016).

(6) P. A. Morris, R. S. Aspden, J. E. Bell, R. W. Boyd, M. J. Padgett, Imaging with a small number of photons. *Nat. Commun.* **6**, 5913 (2015).

(7) A. Gatti, E. Brambilla, M. Bache, L. A. Lugiato, Correlated imaging, quantum and classical. *Phys Rev A* **70**(1), 10 (2004) [doi:10.1103/PhysRevA.70.013802].

(8) M. P. Edgar, G. M. Gibson, R. W. Bowman, B. Sun, N. Radwell, K. J. Mitchell, S. S. Welsh, M. J. Padgett, Simultaneous real-time visible and infrared video with single-pixel detectors. *Sci. Rep*. **5**, 10669–8 (2015) [doi:10.1038/srep10669].

(9) G. A. Howland, D. J. Lum, Photon counting compressive depth mapping. *Opt. Express* **21**(20), 23822–23837 (2013) [doi:10.1364/OE.21.023822].

(10) M.-J. Sun M. P. Edgar, G. M. Gibson, B. Sun, N. Radwell, R. Lamb, M. J. Padgett, Single-pixel three-dimensional imaging with time-based depth resolution. *Nat. Commun.* **7**, 12010–12016 (2016) [doi:10.1038/ncomms12010].

(11) D. B. Phillips, M.-J., Sun, J. M. Taylor, M. P. Edgar, S. M. Barnett, G. M. Gibson, M. J. Padgett, Adaptive foveated single-pixel imaging with dynamic super-sampling. *Sci. Adv.* **3**, e1601782 (2017).

(12) D. V. O'Connor, D. Phillips, *Time-Correlated Single Photon Counting* (UK: Academic Press, London,1984).

(13) W. Becker, *Advanced Time-Correlated Single Photon Counting Techniques* (Springer Series in Chemical Physics, Springer, Berlin, 2005).

(14) S. Pellegrini, G. S. Buller, J. M. Smith, A. M. Wallace, S. Cova, Laser-based distance measurement using picosecond resolution time-correlated single-photon counting. *Meas. Sci. Technol.* **11**(6), 712-716 (2000).

(15) D. Lindell, M. O'Toole, G. Wetzstein, Single-photon 3D imaging with deep sensor fusion. *ACM Trans. Graphics (SIGGRAPH)*, (2018).

(16) A. Kirmani, D. Venkatraman, D. Shin, A. Colaço, F. N. C. Wong, J. H. Shapiro, V. K. Goyal, First-photon imaging. *Science* **343**(6166), 58-61 (2014).

(17) D. Shin, A. Kirmani, V. K. Goyal, J. H. Shapiro, Photon-efficient computational 3D and reflectivity imaging with single-photon detectors. *IEEE Trans. Computat. Imaging* **1**(2), 112-125 (2015).

(18) J. Rapp, V. K. Goyal, A few photons among many: unmixing signal and noise for photon-efficient active imaging. *IEEE Trans. Computat. Imaging* **3**(3), 445-459 (2017).

(19) F. Villa, R. Lussana, D. Bronzi, S. Tisa, A. Tosi, F. Zappa, A. Dalla Mora, D. Contini, D. Durini, S. Weyers, W. Brockherde, CMOS imager with 1024 SPADs and TDCs for single-photon timing and 3-D time-of-flight. *IEEE J. Sel. Topics in Quantum Electronics* **20**(6), 364-373 (2014).

(20) D. Shin, F. Xu, D. Venkatraman, R. Lussana, F. Villa, F. Zappa, V. K. Goyal, F. N. C. Wong, J. H. Shapiro, Photon-efficient imaging with a single-photon camera. *Nat. Commun.* **7**, 12046 (2016).

(21) A. Velten, T. Willwacher, O. Gupta, A. Veeraraghavan, M. G. Bawendi, R. Raskar, Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature Communications* **3**, 745 (2012).

(22) M. Buttafava, J. Zeman, A. Tosi, K. Eliceiri, A. Velten, Non-line-of-sight imaging using a time-gated single photon avalanche diode. *Optics Express* **23**, 20997 (2015).

(23) M. Laurenzis, A. Velten, Non line-of-sight laser gated viewing of scattered photons. *Optical Engineering* **53**(2), 023102 (2014).

(24) J. Klein, C. Peters, J. Martin, M. Laurenzis, M. B. Hullin. Tracking objects outside the line of sight using 2D intensity images. *Sci. Rep.* **6**, 32491 (2016).

(25) G. Gariepy, F. Tonolini, R. Henderson, J. Leach, D. Faccio. Detection and tracking of moving objects hidden from view. *Nat. Photon.* **10**, 23-26 (2016).

(26) S. Chan, R. Warburton, G. Gariepy, J. Leach, D. Faccio. Non-line-of-sight tracking of people at long range. *Opt. Express* **25**, 10109 (2017).

(27) N. Naik, S. Zhao, A. Velten, R. Raskar, K. Bala. Single view reflectance capture using multiplexed scattering and time-of-flight imaging. *ACM Trans. on Graphics* **30**(6), 171 (2011).

(28) R, Pandharkar, A. Velten, A. Bardagjy, E. Lawson, M. Bawendi, R. Raskar. Estimating motion and size of moving non-line-of-sight objects in cluttered environments. *IEEE Conf. Comp. Vision Patt. Recogn. (CVPR),* Providence, RI, USA, 265-272 (2011).

(29) M. O'Toole, D. B. Lindell, G. Wetzstein, Confocal non-line-of-sight imaging with the light cone transform. *Nature* **555**, 338-341 (2018).

(30) A. Kirmani, H. Jeelani, V. Montazerhodjat, V. K. Goyal, Diffuse imaging: creating optical images with unfocused time-resolved illumination and sensing. *IEEE Signal Process. Lett.* **19**(1), 31-34 (2012).

(31) P. Boston *et al.*, http://www.niac.usra.edu/files/studies/final_report/710Boston.pdf (2004)

(32) C. Veerappan, J. Richardson, R. Walker, D.-U. Li, M. W. Fishburn, Y. Maruyama, D. Stoppa, F. Borghetti, M. Gersback, R. K. Henderson, E. Charbon, A 160 x 128 single-photon image sensor with on-pixel 55 ps 10 bit time-to-digital converter. *IEEE Int. Solid-State Circuit Conf. Dig. Tech. Papers (ISSCC)*, 312-314 (2011).

(33) I. Vornicu, R. Carmona-Galan, A. Rodriguez-Vazquez, Arrayable voltage-controlled ring-oscillator for direct time-of-flight image sensors. *IEEE Trans. on Circuits and Systems I: Regular Papers* **64**(11), 2821-

2834 (2017).

(34) L. Gasparini, M. Zarghami, H. Xu, L. Parmesan, M. M. Garcia, M. Unternahrer, B. Bessire, A. Stefanov, D. Stoppa, M. Perenzoni, A 32x32-pixel time-resolved single-photon image sensor with 44.64 μm pitch and 19.48% fill-factor with on-chip row/frame skipping features reaching 800kHz observation rate for quantum physics applications. *IEEE Intern. Solid-State Circuits Conference, San Fransisco, USA* (2018).

(35) G. Gariepy, N. Krstajic, R. Henderson, C. Li, R. R. Thomson, G. S. Buller, B. Heshmat, R. Raskar, J. Leach, D. Faccio. Single-photon sensitive light-in-flight imaging. *Nat. Commun.* **6**, 6021 (2015).

(36) J. A. Giordmaine, P. M. Rentzepis, S. L. Shapiro, K. W. Wecht. Two photon excitation of fluorescence by picosecond light pulses. *Applied Physics Letters* **11**, 216-218 (1967).

(37) M. A. Duguay, J. W. Hansen. An ultrafast light gate. *Applied Physics Letters* **15**, 192-194 (1969).

(38) M. A. Duguay, J. W. Hansen. Ultrahigh-speed photography of picosecond light pulses. *IEEE Journal of Quantum Electronics* **10**(9), 2162-2170 (1971).

(39) A.Velten, D. Wu, A. Jarabo, B. Masia, C. Barsi, C. Joshi, E. Lawson, M. Bawendi, D. Gutierrez, R. Raskar. Femto-photography: capturing and visualizing the propagation of light. *ACM Trans. on Graphics.* **32**, 1-8 (2013).

(40) R. Warburton, C. Aniculaesei. M. Clerici, Y. Altmann, G. Gariepy, R. McCracken, D. Reid, S. McLaughlin, M. Petrovich, J. Hayes, R. Henderson, D. Faccio, J. Leach, Observation of laser pulse propagation in optical fibers with a SPAD camera. *Sci. Rep.* **7**, 43302 (2017).

(41) J.M. Pavia, M. Wolf, E. Charbon, Single-photon avalanche diode images applied to near-infrared imaging. *IEEE J. Sel. Top. Quant. Electron.* **20**, 3800908 (2014).

(42) G. Satat, B. Heshmat, D. Raviv, T Raskar. All photons imaging through volumetric scattering. *Sci. Rep.* **6**, 33946 (2016).

(43) A. Maccarone, A. McCarthy, X. Ren, R. E. Warburton, A. M. Wallace, J. Moffat, Y. Petillot, G. S. Buller, Underwater depth imaging using time-correlated single-photon counting. *Opt. Express* **23**(26), 33911-33926 (2015).

(44) Y. Altmann, X. Ren, A. McCarthy, G. S. Buller, S. McLaughlin, Lidar waveform-based analysis of depth images constructed using sparse single photon data. *IEEE Trans. Image Processing* **25**(5), 1935-1946 (2016).

(45) A. Halimi, A. Maccarone, A. McCarthy, S. McLaughlin, G. S. Buller, Object depth profile and reflectivity restoration from sparse single-photon data acquired in underwater environments. *IEEE Trans. on Computat. Imaging* **3**(3), 472-484 (2017).

(46) P. W. Milonni, J. H. Carter, C. G. Peterson, R. J. Hughes, Effects of propagation through atmospheric turbulence on photon statistics. *J. of Optics B: Quantum and Semiclassical Optics* **6**(S742) (2004).

(47) I. Capraro, A. Tomaello, A. Dall'Arche, F. Gerlin, R. Ursin, G. Vallone, P. Villoresi, Impact of turbulence in long range quantum and classical communications. *Phys. Rev. Lett.* **109**(20), 200502 (2012).

(48) M. Henriksson, L. Sjoqvist, Scintillation index measurement using time-correlated single-photon counting laser radar. *Opt. Eng.* **53**(8), 081902 (2014).

(49) A. McCarthy, X. Ren, A. D. Frera, N. R. Gemmell, N. J. Krichel, C. Scarcella, A. Ruggeri, A. Tosi, G. S. Buller, Kilometer–range depth imaging at 1550 nm wavelength using an InGaAs/InP single-photon avalanche diode detector. *Opt. Express* **21**(19), 22098-22113 (2013).

(50) A. McCarthy, N. J. Krichel, N. R. Gemmell, X. Ren, M. G. Tanner, S. D. Dorenbos, V. Zwiller, R. H. Hadfield, G. S. Buller , Kilometer-range, high resolution depth imaging via 1560 nm wavelength single–photon detection. *Opt. Express* **21**(7), 8904-8914 (2013).

(51) Y. Altmann, X. Ren, A. McCarthy, G. S. Buller, S. McLaughlin, Robust Bayesian target detection algorithm for depth imaging from sparse single-photon data. *IEEE Trans. on Computat. Imaging* **2**(4), 456-

467 (2016).

(52) A. M. Pawlikowska, A. Halimi, R. A. Lamb, G. S. Buller, Single-photon three-dimensional imaging at up to 10 kilometers range. *Opt. Express* **25**(10), 11919-11931 (2017).

(53) J. Zhu, C. Yajun, Z. Labao, J. Xiaoqing, F. Zhijun, W. Ganhua, Y. Xiachao, J. Zhai, Y. Wu, Q. Chen, X. Zhou, Z. Wang, C. Zhang, L. Kang, J. Chen, P. Wu, Demonstration of measuring sea fog with an SNSPD-based Lidar system. *Sci. Rep.* **7**, 15113 (2017).

(54) B. Du, C. Pang, D. Wu, Z. Li, H. Peng, Y. Tao, E. Wu, G. Wu, High-speed photon-counting laser ranging for broad range of distances. *Sci. Rep.* **8**, 4198 (2018).

(55) G. Satat, B. Heshmat, D. Raviv, R. Raskar, All photons imaging through volumetric scattering. *Sci. Rep.* **6**, 33946 (2016).

(56) G. Satat, M. Tancik, R. Raskar, Towards photography through realistic fog. *IEEE Intern. Conf. on Computat. Photography (ICCP),* Pittsburgh, PA, USA (2018).

(57) F. H. Goetz, Three decades of hyperspectral remote sensing of the Earth: a personal view. *Remote Sensing of Environment* **113**(S1), S5-S16 (2009).

(58) A. A. Gowen, C. P. O'Donnell, P. J. Cullen, G. Downey, J. M. Frias, Hyperspectral imaging - an emerging process analytical tool for food quality and safety control. *Trends in Food Science & Technology* **18**(12), 590-598 (2007).

(59) G. P. Asner, D. E. Knapp, T. Kennedy-Bowdoin, M. O. Jones, R. E. Martin, J. W. Boardman, C. B. Field, Carnegie airborne observatory: in-flight fusion of hyperspectral imaging and waveform light detection and ranging for three-dimensional studies of ecosystems. *J. of Applied Rem. Sens.* 1(1), 013536 (2007).

(60) D. Landgrebe, Hyperspectral image data analysis. *IEEE Sig. Process. Mag.* **19**(1), 17-28 (2002).

(61) A. Plaza, J. A. Benediktsson, J. W. Boardman, J. Brazile, L. Bruzzone, G. Camps-Valls, J. Chanussot, M. Fauvel, P. Gamba, A. Gualtieri, M. Marconcini, Recent advances in techniques for hyperspectral image processing. *Remote sensing of environment* **113**(S1), S110-S122 (2007).

(62) J. M. Bioucas-Dias, A. Plaza, N. Dobigeon, M. Parente, Q. Du, P. Gader, J. Chanussot, Hyperspectral unmixing vverview: geometrical, statistical, and sparse regression-based approaches. *IEEE J. of Sel. Topics in Appl. Earth Observ. and Rem. Sens*. **5**(2), 354-379 (2012)

(63) M. Fauvel, Y. Tarabalka, J. A. Benediktsson, J. Chanussot, J. C. Tilton, Advances in spectral-spatial classification of hyperspectral images. *Proc. of the IEEE* **101**(3), 652-675 (2013).

(64) N. Dobigeon, J.-Y. Tourneret, C. Richard, J. C. M. Bermudez, S. McLaughlin, A. O. Hero, Nonlinear unmixing of hyperspectral Images: models and algorithms. *IEEE Sig. Process. Mag.* **31**(1), 82-94 (2014).

(65) L. Loncan, L. B. Almeida, J. M. Bioucas-Dias, X. Briottet, J. Chanussot, N. Dobigeon, S. Fabre, W. Liao, G. A. Licciardi, M. Simões, J.-Y. Tourneret, M. A. Veganzones, G. Vivone, Q. Wei, N. Yokoya, Hyperspectral pansharpening: a review. *IEEE Geosc.and Rem. Sens. Mag.* **3**(3), 27-46 (2015).

(66) D.-W. Sun, *Hyperspectral Imaging for Food Quality Analysis and Control* (Academic Press, 2010)

(67) R. Koprowski, O. Paweł, Segmentation in dermatological hyperspectral images: dedicated methods. *Biomed. Eng. Online* **15**(1), 97 (2016).

(68) F. Vasefi, N. MacKinnon, D. L. Farkas, "Chapter 16 - Hyperspectral and Multispectral Imaging in Dermatology" in *Imaging in Dermatology* (Academic Press, Boston, 2016).

(69) A. M. Wallace, A. McCarthy, C. J. Nichol, X. Ren, S. Morak, D. Martinez-Ramirez, I. H. Woodhouse, G. S. Buller, Design and Evaluation of Multispectral Lidar for the recovery of arboreal parameters. *IEEE Trans. Geosci. and Rem. Sensing* **52**(8), 4942-4954 (2014).

(70) Y. Altmann, A. Maccarone, A. McCarthy, G. Newstadt, G. S. Buller, S. McLaughlin, A. Hero, Robust spectral unmixing of sparse multispectral Lidar waveforms using gamma Markov random fields. *IEEE Trans. on Computat. Imaging* **2**(4), 658-670 (2017).

(71) P. Chhabra, A. Maccarone, A. McCarthy, G. S. Buller, A. Wallace, Discriminating underwater Lidar target signatures using sparse multi-spectral depth codes. *Sensor Signal Processing for Defence (SSPD)*,

Edinburgh, U.K. (2016).

(72) Y. Altmann, A. Maccarone, A. McCarthy, G. S. Buller, S. Mclaughlin, Joint spectral clustering and range estimation for 3D scene reconstruction using multispectral Lidar waveforms. *Europ. Signal Process. Conf. (EUSIPCO)*, Budapest, Hungary, 513-517 (2016).

(73) Y. Altmann, A. Maccarone, A. McCarthy, G. Buller, S. McLaughlin, Joint range estimation and spectral classification for 3D scene reconstruction using multispectral Lidar waveforms. *IEEE Statist. Signal Process. Work. (SSP),* Palma de Mallorca, Spain, 1-5 (2016).

(74) Y. Altmann, A. Maccarone, A. McCarthy, S. McLaughlin, G. S. Buller, Spectral classification of sparse photon depth images. *Opt. Express* **26**(5), 5514-5530 (2018).

(75) R. Tobin, Y. Altmann, X. Ren, A. McCarthy, R. A Lamb, S. McLaughlin, G. S. Buller, Comparative study of sampling strategies for sparse photon multispectral lidar imaging: towards mosaic filter arrays. *J. of Opt.* **19**(19), 094006 (2017).

(76) Y. Altmann, R. Tobin, A. Maccarone, X. Ren, A. McCarthy, G. S. Buller, S. McLaughlin, Bayesian restoration of reflectivity and range profiles from subsampled single-photon multispectral Lidar data. *Europ. Signal Process. Conf. (EUSIPCO)*, Kos, Greece, 1410-1414 (2017).

(77) R. Takagi, R. Horisaki, J. Tanida, Object recognition through a multi-mode fiber. *Opt. Rev.* **24**(2), 117-120 (2017).

(78) A. Sinhai, J. Leee, S. Li, G. Barbastathis, Lensless computational imaging through deep learning. *Optica* **4**(9), 1117-1125 (2017).

(79) G. Satat, M. Tancik, O. Gupta, B. Heshamt, R. Raskar, Object classification through scattering media with deep learning on time resolved measurement. *Opt. Express* **25**(15), 17466-17479 (2017).

(80) P. Caramazza *et al.*, https://arxiv.org/abs/1709.07244 (2017).

(81) D. Shin, F.Xu, F. N. C. Wong, J. H. Shapiro, V. K Goyal, Computational multi-depth single-photon imaging. *Opt. Express* **24**(3), 1873-1888 (2016).

(82) R. Tobin, A. Halimi, A. McCarthy, X. Ren, K. J. McEwan, S. McLaughlin, G. S. Buller, Long-range depth profiling of camouflaged targets using single-photon detection. *Opt. Eng.* **57**(3) 031303 (2017).

(83) S. Medin, J. Murray-Bruce, V. K. Goyal, Optimal stopping times for estimating Bernoulli parameters with applications to active imaging. *IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP),* Calgary, Canada (2018).

Acknowledgments: The authors thank the Royal Society for support and hosting the Theo Murphy Scientific meeting on "Light Transport and imaging through complex media.

Funding: DF acknowledges financial support from the UK Engineering and Physical Sciences Research Council (grants EP/ M006514/1, EP/ M01326X/1). YA acknowledges the support from the UK Royal academy of Engineering under the Research Fellowship Scheme (RF201617/16/31). VKG acknowledges support from the US Defence Advanced Research Projects Agency (DARPA) InPho program through U.S. Army Research Office award W911NF-10-1-0404, the US DARPA REVEAL program through contract HR0011-16-C-0030, and US National Science Foundation through grants 1161413 and 1422034. AOH acknowledges the support from the U.S. Army Research Office award W911NF-15-1-0479 and the U.S. Dept. of the Air Force grant FA8650-15-D-1845.

Author contributions: all authors contributed equally to this work.

Data and materials availability: No new data was generated in this work.

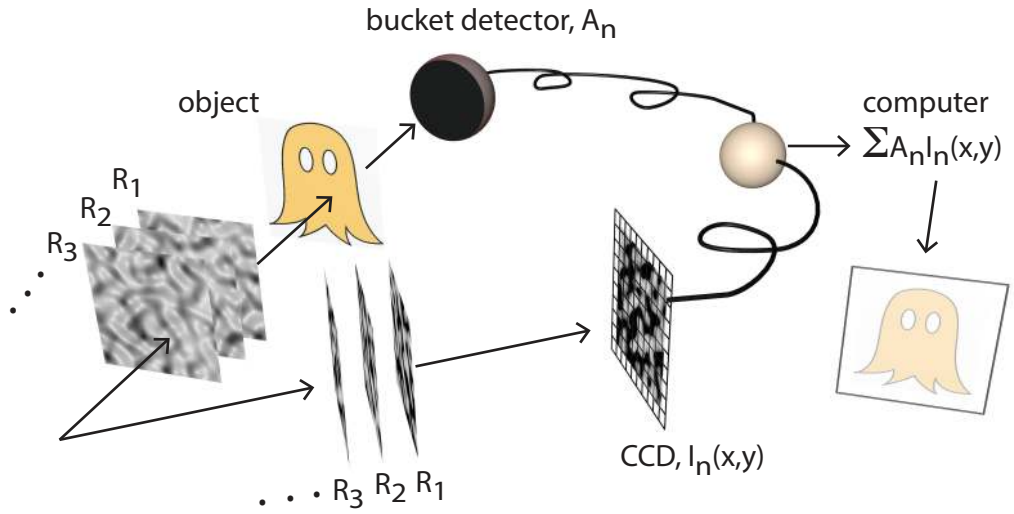

Fig.1: Ghost imaging. Random spatial patterns, R_n, illuminate an object and only the total transmitted (or reflected) light, A_n, is measured. This intensity reading is then computationally combined with the information of the pattern, $I_n(x,y)$ (either measured separately or known if this was generated by a computer) to form an image of the object.

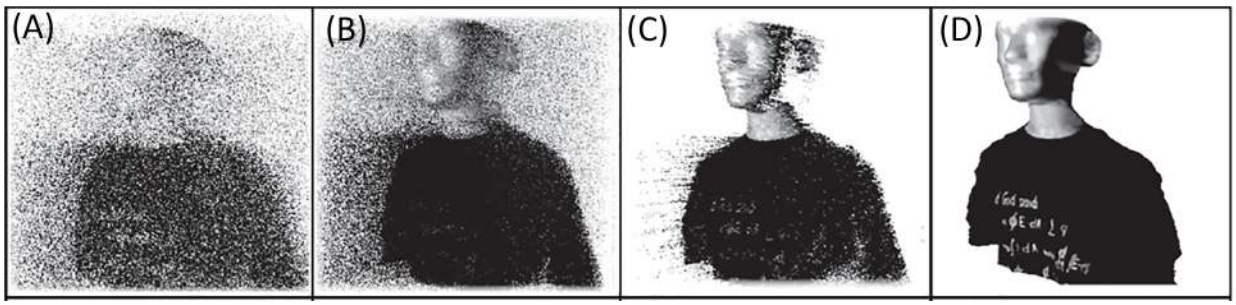

Fig.2 First-photon imaging. (A) Each photon detection can be mapped to a 3D position, which is often far from correct because half of the detections are due to background light. The number of illumination pulses until a photon is detected is inversely proportional to an initial estimate of reflectivity. (B) Exploiting piecewise smoothness yields improved reflectivity estimates. (C) Approximate noise censoring removes most detections due to background light. (D) The final estimate exploits piecewise smoothness of depth. Figure adapted from [*(15)*, Fig. 2]

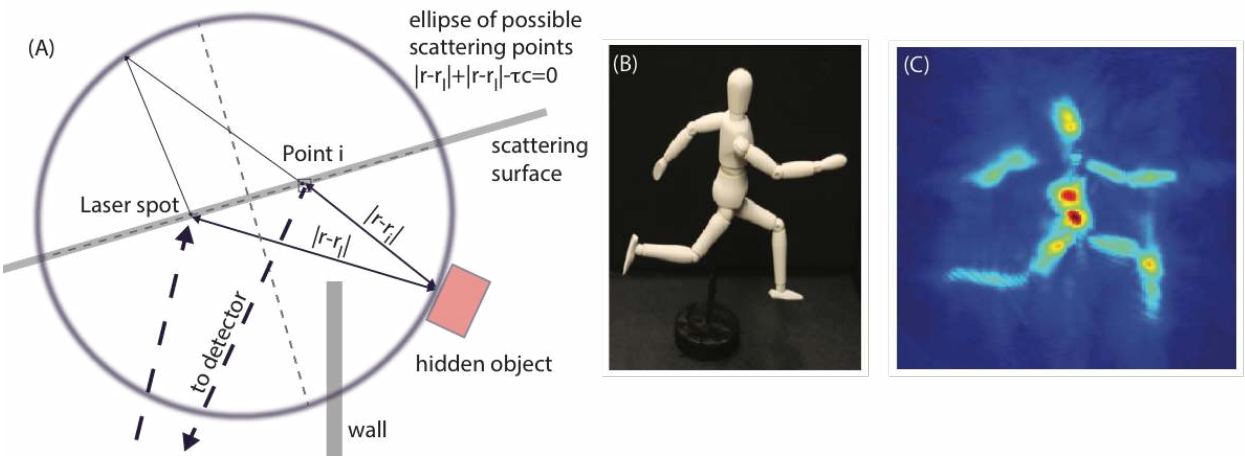

Fig.3 Non-Line-Of-Sight imaging. (A) Basic geometry of the problem. A laser illuminates a scattering surface and scatters light around a wall that hides an object from the direct line of sight. The return signal back-scattered from the hidden object is detected at a point "i" on the scattering surface. This geometry, with a single observation point, defines an ellipsoid of possible locations for the object. Detection of the time-resolved transient image at multiple points on the surface allows to reconstruct the location or even the full 3D shape of the object. (B) An example of a hidden object with its reconstruction shown in (C). Figures (B) and (C) adapted from *(20)*.

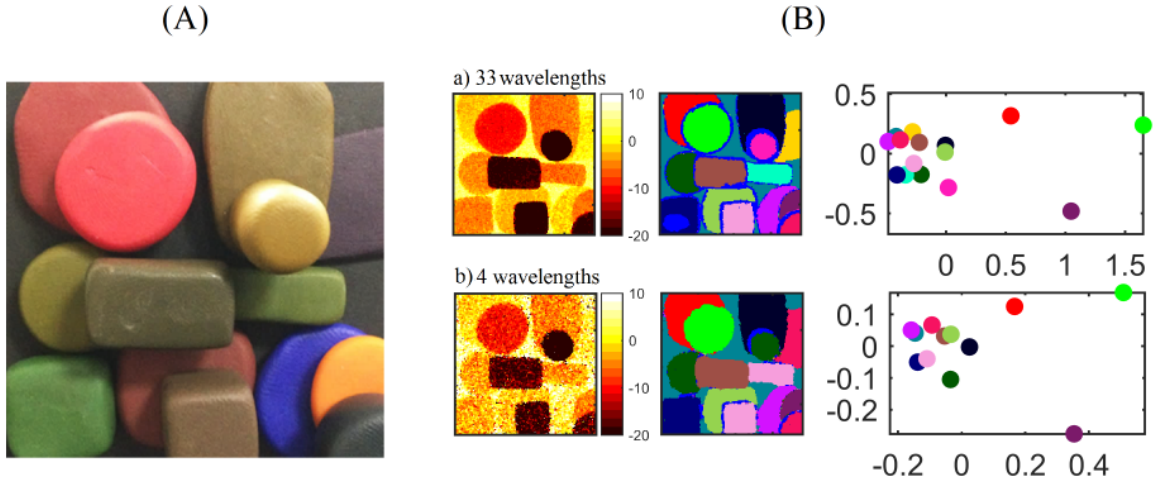

Fig.4 Computational inverse probability methods to spectrally classify and depth resolve objects in a scene from photon-starved multispectral LIDAR images. The scene (A) was composed of 14 clay materials composed of different colours. The recorded images consist of 200 × 200 pixels (the scanned target areas were approximately 50 × 50 mm) and the targets were placed at 1.85 meters from the system. In (B), the first column depicts the estimated depth profile (in mm), the reference range being arbitrarily set to the backplane range. The second column depicts color classification maps and the third column depicts the spectral signatures of the most prominent classes, projected onto the first and second axes obtained using principal component analysis. Each of these subplots illustrates the similarity between the estimated spectral signatures. Rows a) and b) in (B) depict

results obtained with an average of 1 detected photon per pixel, for each spectral band, with 33 and only 4 bands. Figure adapted from *(63)*