

Region-of-Interest Prioritised Sampling for Constrained Autonomous Exploration Systems

Protim Bhattacharjee , Martin Burger , Anko Börner, and Veniamin I. Morgenshtern 

Abstract—Goal oriented autonomous operation of space rovers has been known to increase scientific output of a mission. In this work we present an algorithm, called the RoI Prioritised Sampling (RPS), that prioritises Region-of-Interests (RoIs) in an exploration scenario in order to utilise the limited resources of the imaging instrument on the rover effectively. This prioritisation is based on an estimator that evaluates the change in information content at consecutive spatial scales of the RoIs without calculating the finer scale reconstruction. The estimator, called the Refinement Indicator (RI), is motivated and derived. Multi-scale acquisition approaches, based on classical and multi-level compressed sensing, with respect to the single pixel camera architecture are discussed. The performance of the algorithm is verified on remote sensing images and compared with the state-of-the-art multi-resolution reconstruction algorithms. At the considered sub-sampling rates the RPS is shown to better utilise the system resources for reconstructing the RoIs.

Index Terms—Autonomous systems, compressed sensing, image acquisition, space exploration, spatial resolution.

I. INTRODUCTION

IMAGE acquisition in resource constrained environments is a challenging task. Activities such as space exploration and investigation of disaster sites are carried out by robotic platforms that have limited electrical power and payload capabilities. Being far away from the ground-station/human operator, exploration platforms, e.g., rovers and copters, require autonomous operation for achieving mission objectives. This necessitates new on-board sensing protocols, object and event detection algorithms, and data processing frameworks [1], [2]. The imaging systems on autonomous exploration platforms are required to acquire maximum information about the scene with minimum resources. Limitations on electrical power impact the

number of measurements that can be acquired and the number of computations that can be performed by the imaging system. For such constrained systems efficient utilization of their limited resources is important. The main purpose of the exploring rovers is to provide a survey of the scene, so as to recognise areas of importance, and possibly, to provide initial data for further examination by more precise instruments. This reduces the requirement of acquiring the entire scene at the highest resolution of the camera on-board the robotic platform. Regions-of-Interest (RoIs) can be recognized and acquired at the resolution of the camera leaving the background at lower resolutions.

To ensure that measurements are expended on RoIs with more information it is necessary to prioritise RoIs for acquisition on the basis of their information content. Moreover, this should be performed in an online manner. However, existing methods of acquisition and reconstruction in constrained systems either try to reconstruct the entire scene at low resolutions [3], [4] or require pre-defined resolutions for each RoI [5]. These methods do not dynamically distribute their measurements among RoIs based on their information content. Also, they may require prior information (human intervention or some other source) to decide on the spatial resolution of each RoI.

Dynamic real-time multi-resolution RoI acquisition based on the information content of the RoI is an important step forward for introducing mission oriented autonomy to exploration systems (E4 level autonomy) [6]. There are various ways of characterising the information content of a scene or RoI. In this work an adaptive multi-scale approach is used. An estimator, called the Refinement Indicator (RI), is developed to estimate changes in the information content of a RoI at consecutive spatial resolutions. The RoI with the largest change is refined to a finer spatial resolution. With the help of the RI and the basic framework for sampling suggested in [7] for measurement-constrained systems an algorithm for multi-scale sampling is proposed and named RoI Prioritised Sampling (RPS). The RI provides a structured way to prioritise RoIs for acquisition at the cost of a small number of overhead measurements at each spatial resolution. These additional measurements are also used for reconstruction of the RoIs using compressed sensing [8], [9].

The main contributions of this work are 1) the development of an online estimator for information change across spatial scales for RoIs, 2) the design of an acquisition algorithm for constrained systems that incorporates this estimator to prioritise RoIs for acquisition and accordingly distributes the limited measurement budget, and 3) the proposal of multi-scale acquisition

Manuscript received April 6, 2021; revised September 17, 2021 and January 24, 2022; accepted March 12, 2022. Date of publication March 30, 2022; date of current version April 18, 2022. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Jinwei Gu. (Corresponding author: Protim Bhattacharjee.)

Protim Bhattacharjee and Anko Börner are with the Department of Optical Sensor Systems, German Aerospace Center (DLR), 12489 Berlin, Germany (e-mail: protim.bhattacharjee@dlr.de; anko.boerner@dlr.de).

Martin Burger is with the Department of Mathematics, Friedrich Alexander Universität Erlangen-Nürnberg, 91054 Erlangen, Germany (e-mail: martin.burger@fau.de).

Veniamin I. Morgenshtern is with the Department of Electrical-Electronic-Communication Engineering, Friedrich Alexander Universität Erlangen-Nürnberg, 91054 Erlangen, Germany (e-mail: veniamin.morgenshtern@fau.de).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TCI.2022.3163552>, provided by the authors.

Digital Object Identifier 10.1109/TCI.2022.3163552

methodologies based on classical [10], [11] and multi-level compressed sensing [11]–[13]. The algorithm is tested on images from airborne Earth observation sensor platforms and the Mars Science Laboratory (MSL) on-board the Curiosity rover on Mars. The proposed algorithm is found to perform better at reconstructing informative RoIs than existing algorithms at the considered sub-sampling rates.

The next section reviews existing work on autonomous space exploration and compressed sensing based acquisition in constrained systems. Section III discusses the camera model used in this work and the sampling methodology suggested in [7]. The RI is motivated and derived in Section IV. Section V discusses the datasets and results as applied to these datasets. The article ends with a conclusion and the direction for future work.

II. LITERATURE REVIEW

A. Autonomous Space Exploration

Autonomous Exploration for Gathering Increased Science (AEGIS) [2] was developed for the Opportunity rover [14] on Mars to increase the scientific output of the mission by allowing autonomous acquisition of mission relevant data. The aim of the AEGIS system is to detect rocks by analysing images from the navigation camera of the rover and to direct scientific instruments, like the ChemCam [15], to acquire scientific data from the surroundings of the rover without consultation with Earth scientists. Earlier all images were downlinked and analysed on Earth and then the rovers were instructed to perform extra measurements or revisit an area from which they had moved away. However, with the AEGIS system new scientific information could be gathered autonomously by the rover leading to increased scientific throughput [16]. The current work builds on the idea of autonomous scientific data collection and provides a method for prioritising (ranking) the RoIs based on spatial frequency variations without any prior knowledge about the surrounding terrain.

B. Compressed Sensing for Constrained Systems

Multi-resolution techniques have previously been used to address concerns of constrained systems. In [3] the aim is to provide low resolution previews in computationally-constrained and data-streaming systems where computational machinery for compressed sensing reconstruction is not available. The authors propose an orthogonal sampling matrix known as the Sum-To-One (STOne) transform. This matrix allows one to recover low resolution previews at the Nyquist-rate in real-time. The sampling matrix contains only ± 1 entries and is suitable for implementation on a bi-stable Digital Micromirror Device (DMD) used in Single Pixel Cameras (SPCs). The nature of sampling is such that when computational resources are available (at the ground/base-station) the measurements used for the preview can also be used to generate a high resolution image at sub-Nyquist rates. The previews can be generated at any resolution that is a power of 2 limited by the size of the DMD. The STOne transform thus tries to acquire and reconstruct the entire scene in a computationally constrained system. This scenario can also be thought

of as a measurement-constrained problem where the limited measurements are designed such that the exploration system can produce low resolution images for autonomous functions and high resolution recovery is possible only at a ground-station. Authors in [4] cast the measurement-constrained reconstruction problem into a Multi-Resolution Approximate Message Passing (MR-AMP) framework. The measurement model is transformed in such a way that the low resolution image is an optimization variable. Properly designed down-sampling and up-sampling matrices are required to map the image from the native resolution of the sensor to the lower resolution and vice-versa. A down-sampling factor can be used to choose the down-sampling ratio according to the available number of measurements. As in the case of the STOne transform, MR-AMP tries to reconstruct the complete scene at the desired resolution. Multi-resolution in RoIs in the context of compressed sensing camera architectures was proposed in [5]. Measurements for the entire scene are acquired at the native resolution of the DMD. The original image is then split into a number of pre-defined RoIs, and with the help of down-sampling and up-sampling matrices, different RoIs are reconstructed at different resolutions. The RoIs need to be chosen off-line and are reconstructed from the same measurements. The three techniques described model the reconstruction problem in a manner that enables multi-resolution recovery. In [7], the authors propose a different acquisition approach for measurement-constrained systems. A three-step methodology for acquiring RoIs in a scene is proposed that includes a “Low resolution acquisition and reconstruction” step, “RoI detection and segmentation” step, and a “Multi-level sampling” step. The entire process is online; it is performed in-situ without any external guidance or human intervention. The segmentation procedure is performed on the basis of mission objectives and multi-level compressed sensing is used to recover the RoIs. Each RoI is assigned a measurement budget on the basis of its size and is acquired individually. However, the RoI selection procedure is empirical and the algorithm acquires RoIs based on their sizes and tries to reconstruct each RoI at the native resolution of the DMD in one shot with no regard to the information content of the RoI. We propose to estimate the change in information content across spatial scales in the RoIs in an online fashion and distribute the measurement budget accordingly. A step-by-step increase in resolution is obtained in the RoIs by acquiring them at different resolutions instead of reconstructing them directly from measurements at the native resolution of the DMD.

C. Bregman Distance

In this work, Bregman distance [17] will be used to calculate error estimates. For a convex functional J with a subgradient $p \in \partial J(x)$, the (generalized) Bregman distance between two vectors z and x is defined as $D_J^p(z, x) = J(z) - J(x) - \langle p, z - x \rangle$, where $\partial J(x)$ is the subdifferential of J at x . The Bregman distance is the distance between $J(z)$ and the tangent to J at x evaluated at z . This is shown in Fig. 1. For a convex functional J , the non-negativity is evident from the definition of J . However, the Bregman distance is not necessarily symmetric; a symmetric Bregman distance [17] with respect

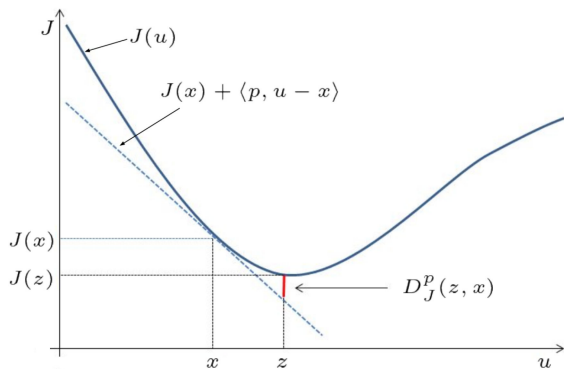


Fig. 1. Bregman Distance is the distance, displayed in red, between $J(z)$ and the tangent to $J(\cdot)$ at x evaluated at z .

to two subgradients of J , $p \in \partial J(z)$ and $q \in \partial J(x)$, is defined as $D_J^{p,q}(z, x) = D_J^p(x, z) + D_J^q(z, x) = \langle p - q, z - x \rangle$. This symmetric form of the Bregman distance will be used to estimate the error in reconstruction of RoIs at various spatial scales. Similar methods for deriving error estimates for regularisation problems [18] and image restoration problems [19] based on Bregman distances have been studied previously. For l_1 regularised problem, common in compressed sensing, the Bregman distance is related to the sparsity of the signal [19, Proposition 8.2].

III. SINGLE PIXEL CAMERA MODEL AND ACQUISITION METHODOLOGY

A. Single Pixel Camera

Different applications require acquisition of the scene at different wavelengths of the electromagnetic spectrum. For example, fire or temperature detection of objects require near-to-mid infrared imaging [20]. Terahertz radiation is used in security applications [21]. Building large sensor arrays for imaging in longer wavelengths is difficult and expensive. The alternative is to use a single detector appropriate for the incident wavelength. Such a detector when calibrated is known as a radiometer [22]–[24]. A mechanical system can be used to raster scan the scene with the radiometer to generate a 2D image of the scene. However, in space applications mechanical systems are susceptible to damage due to stress and vibrations during rocket launch and landing. SPCs [25] are best suited for such exploration scenarios as they avoid large sensor arrays and mechanical scanning. In this work, the SPC architecture is used as the model for the imaging system. It consists of a camera lens that focusses the incoming radiation onto the spatial light modulator that is placed at the virtual imaging plane of the camera lens. The DMD [26] is used as the spatial light modulator. It provides the necessary “software” scan of the scene replacing mechanical raster scanning. The modulated radiation from the DMD may be filtered through a wavelength selective filter or a colour filter. The collective optics focusses the modulated filtered radiation onto the single pixel detector. At each measurement step, the DMD implements a measurement mask and the incoming radiation

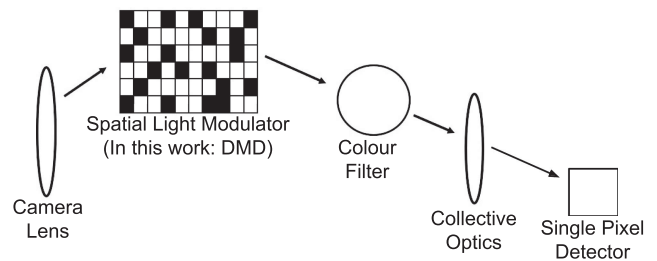


Fig. 2. Single Pixel Camera model.

is encoded with the mask. This encoded radiation is received by the detector and generates one measurement. A sequence of such masks is displayed on the DMD with each mask generating a new measurement. A bi-stable DMD where the micromirrors can only have an ‘ON’ or ‘OFF’ state is used. By convention, the ‘ON’ state directs the incident radiation towards the detector and the ‘OFF’ state deflects radiation away from the detector. For a measurement mask consisting of 0/1 entries, the “0” pixels are mapped to the ‘OFF’ state and the “1” to the ‘ON’ state. Such masks are used to perform random-macro-pixel sampling as discussed in the next subsection. To implement ± 1 entries of a measurement mask, two physical measurement cycles are required on the DMD. The first cycle maps all +1s to the ‘ON’ state and the -1 s to the ‘OFF’ state. In the next measurement cycle all the +1s are mapped to the ‘OFF’ state and the -1 s to the ‘ON’ state. Subtracting the two measurements thus obtained leads to one measurement generated by the measurement mask with ± 1 entries. Walsh transforms [27] require such implementations on the DMD. The basic single pixel camera model is shown in Fig. 2.

The scene is acquired using the three step acquisition methodology proposed in [7]. The first step is the low resolution acquisition that is performed using macro pixels formed by binning micromirrors of the DMD. A 8×8 macro pixel is formed by binning 64 micromirrors underlying a 8×8 window on the DMD and assigning them the same ‘ON’ or ‘OFF’ state. A 256×256 DMD has 1024 non-overlapping 8×8 macro pixels. Similarly, it has 4096 non-overlapping 4×4 macro pixels. For each measurement, these macro pixels are randomly assigned the value of 0/1 or ± 1 to simulate random-macro-pixel sampling. Details on the low resolution reconstruction will follow in Section III-D. The second step is the detection and segmentation based on the low resolution reconstruction. This process depends on the mission objectives. For example, in infrared imaging, one may use temperature as a parameter for detection and segmentation of RoIs; in visible wavelength imaging, one may use the contrast of the scene. The process is same as in [7]. First, the low resolution reconstruction is resized to its macro pixel image dimensions. For example, a 256×256 image with a macro pixel size of 8×8 would be resized to a 32×32 image by averaging each 8×8 block in the low resolution reconstruction. A seed pixel is assigned to be the brightest pixel in the resized image. A square region is grown around the seed pixel by increasing the region size by one pixel in each direction until the contrast of the region is higher than a user defined threshold. This forms one RoI.

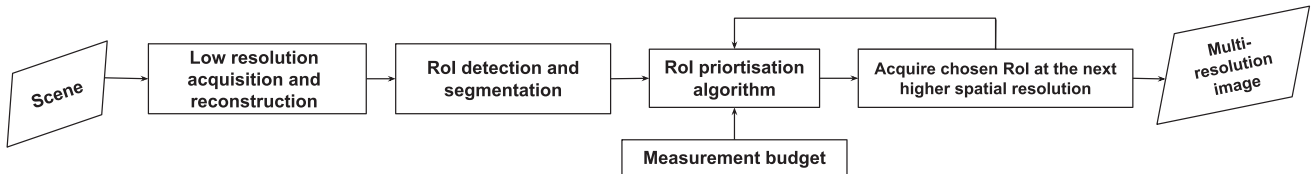


Fig. 3. Proposed acquisition methodology.

A new seed pixel is assigned to be the brightest pixel outside the already selected region. Regions within a merge radius of each other are coalesced forming a larger RoI. This seed pixel selection and region growing is continued until the user defined total number of RoIs is reached. Once all the RoIs are selected, the segmented image is resized to the original size of the low resolution reconstruction, equal to the DMD size. Thereafter, each RoI is expanded to its immediate higher dyadic size. If this leads to more than 99% overlap between RoIs, they are merged once more. The total number of regions to be selected, the contrast cut-off, the merge radius, and the overlap in the dyadic sizes are user defined. Fig. 1 in the supplementary material describes the process. The third step of the algorithm in [7] is the sequential acquisition of RoIs detected in the second step at the full resolution of the DMD using multi-level sampling [12]. The RoIs are arranged in a decreasing manner according to their size and are resolved to their full resolution in a one-shot manner. In this work we propose to estimate the change of information content in spatial scales of the RoIs and develop a procedure for step-by-step increase in the resolution of the RoIs. We propose two ways in which this step-by-step increase can be performed. One is through multi-scale random-macro-pixel sampling and the other is through the Walsh transform. The RoIs are acquired sequentially based on their information content and the number of measurements needed at each spatial scale depends on the RoI size and the current spatial resolution. The workflow of the proposed acquisition methodology is shown in Fig. 3. We first discuss the methods of acquisition of higher spatial frequencies as it will be required for the development of the estimator.

B. Multi-Scale Random-Macro-Pixel Sampling

The process of binning DMD micromirrors to form macro pixels was discussed in the previous section. Each macro pixel sums the incoming radiation from its field of view and the reciprocal of the size of the macro pixel acts as a cut-off for the spatial frequencies acquired by the measurement masks of that macro pixel size. By reducing the macro pixel size one acquires higher spatial frequencies. The macro pixel size of 1×1 defines the native resolution of the DMD. Examples of DMD mirrors binned by different macro pixel size are shown in Fig. 4. This method of multi-scale random-macro-pixel sampling will be used to acquire RoIs at different spatial resolutions.

C. Multi-Scale Multi-Level Sampling With Walsh Transform

An alternate way to sample in a multi-scale manner is to use structured matrices like the Walsh transform. Structured matrices have been used in compressed sensing acquisition

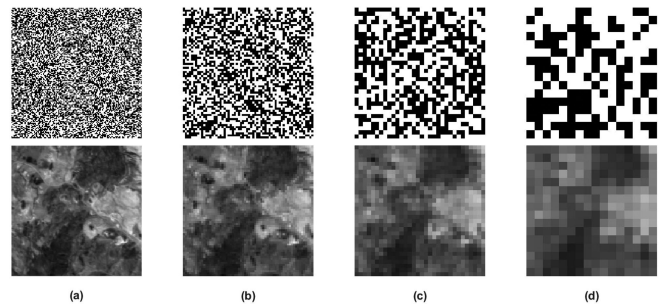


Fig. 4. Illustration of macro pixels. *Top panel*: Example of macro pixel measurement masks formed by binning DMD micromirrors. The macro pixels are randomly assigned 0 or 1 values. *Bottom panel*: Macro pixel representation of the Cuprite Vis dataset obtained by replacing each macro pixel block with the mean of the corresponding block from the original image. Macro pixel size (a) 1×1 (Original), (b) 2×2 , (c) 4×4 , and (d) 8×8 .

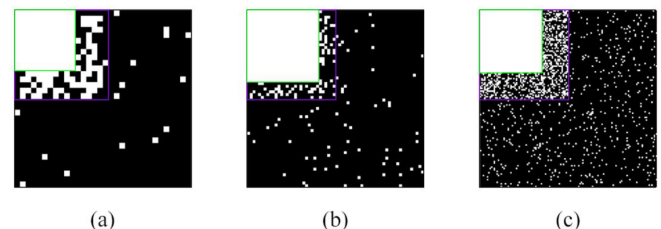


Fig. 5. Walsh transform multi-level sampling maps for RoI of size (a) 32×32 , (b) 64×64 , and (c) 128×128 . Top left corner denotes the DC frequency. Number of measurements for each size is 20% of the total number of pixels in the RoI. The Walsh frequency grids are scaled to the same size for visualisation.

as they can be implemented on bi-stable DMDs, and have fast forward and inverse transform implementations [11], [28]. Another advantage of using Walsh transforms and other structured matrices is that their multi-scale decomposition can be adapted to best leverage the sparsity structure of the signal to be acquired. Multi-level sampling was suggested in [11] and was further developed in [12], [13], [29]. In multi-level sampling the frequency space of the measurement matrix is divided into a number of levels or regions based on its coherence structure and the number of measurements are assigned to these levels in an asymptotically decreasing manner. The Transform Point Spread Function (TPSF) [29] is used as a metric to generate sampling maps for different sized RoIs. In this work, square regions (as shown in Fig. 5) were chosen instead of the regular circular regions as they allowed easier decomposition of the frequency space into multiple levels. The Walsh frequency grid for each RoI is divided into three levels corresponding to low, mid, and high spatial frequencies. To perform multi-scale sampling we

start from the low frequency region followed by level-wise spatial Walsh frequency measurements to achieve a step-by-step increase in resolution. The term multi-scale is a misnomer for multi-level sampling as multi-level sampling itself is defined in a multi-scale manner; however, we use it in the current context to make the stepwise multi-scale acquisition of the scene explicit. Sampling maps for three different size of RoIs are shown in Fig. 5. The spatial frequency regions may be divided into more than three regions depending upon the application at hand. More information about the number of measurements performed in each spatial frequency region is provided in Section V.

D. Compressed Sensing Algorithms

To complete the discussion on acquisition methodology the compressed sensing reconstruction algorithms used are discussed. Two algorithms are used, the Analysis BPDN [30], [31] and the Analysis+TV [29], [30]. The Analysis BPDN can be written as

$$\underset{x}{\text{minimize}} \|Wx\|_1 + \gamma\|y - Ax\|_2^2, \quad (1)$$

where A is the measurement matrix with normalized columns, W is the sparsity basis, y is the vector of measurements, x is the signal to be recovered, and γ is the regularisation parameter. For a signal $x = [x_1, x_2, \dots, x_N]^T$ of length N , the l_1 norm is defined as $\|x\|_1 = \sum_{i \in N} |x_i|$ and the l_2 norm is defined as $\|x\|_2 = \sqrt{\sum_{i \in N} |x_i|^2}$. Algorithm (1) is used for reconstruction of the low resolution image in the first step of the acquisition process with 2D-DCT as the sparsity basis.

The Analysis+TV algorithm

$$\begin{aligned} &\underset{x}{\text{minimize}} \beta_1 \|Wx\|_1 + \beta_2 \|x\|_{TV} \\ &\text{subject to } \|y - Ax\|_2^2 \leq \eta, \end{aligned} \quad (2)$$

adds a Total Variation (TV) regularisation term to the Analysis BPDN. The TV norm for images is calculated along the horizontal and vertical directions [32], $\|x\|_{TV} = \sum_{i,j \in N} \sqrt{|x_{i+1,j} - x_{i,j}|^2 + |x_{i,j+1} - x_{i,j}|^2}$. Such a model is used when the smoothness, promoted by the TV norm, and sparsity constraints, promoted by the W basis, are required simultaneously. The terms β_1 and β_2 are used to balance the contributions of the two prior terms, and η is used to control the size of the noise.

In this section the camera model and the methods for step-by-step increase in the acquisition of spatial frequencies were discussed. The next section motivates and derives the information change estimator, which is used to develop the proposed RoI prioritisation and acquisition algorithm.

IV. ROI PRIORITISED SAMPLING FOR CONSTRAINED SYSTEMS

A. Refinement Indicator

After the ‘‘RoI detection and segmentation’’ step in the proposed acquisition scheme shown in Fig. 3, we must distribute the limited measurement budget among the RoIs based on their information content or, in other words, prioritise the RoIs for sequential acquisition. Let $x \in \mathbf{R}^N$ be the RoI to be acquired at

the finest resolution. Let us start by acquiring some coarse measurements. The system of equations for the coarse measurements is

$$y^C = A^C x + \eta^C, \quad (3)$$

where $A^C = (a_i^C)_{i=1,\dots,m} \in \mathbf{R}^N$ is the coarse resolution measurement matrix, the coarse measurements are y^C , and η^C is an additive noise component. Denoting the solution to the regularised problem as x^C , we have

$$x^C = \arg \min_z \frac{1}{2} \|A^C z - y^C\|_2^2 + \alpha J(z), \quad (4)$$

where J is a regularisation functional like total variation or a wavelet norm and $\alpha \geq 0$ is a regularisation parameter. The optimality condition for (4) [33], [34] is given by

$$(A^C)^*(A^C x^C - y^C) + \alpha p^C = 0, \quad p^C \in \partial J(x^C), \quad (5)$$

where p^C is a subgradient [34] of J at x^C and $\partial J(x^C)$ denotes the subdifferential [34] of J at x^C . The adjoint of A^C is denoted by $(A^C)^*$. Let us ignore the noise term for the moment so that $y^C = A^C x$. To improve the resolution we acquire extra measurements with a refinement matrix B . This matrix is composed of sampling patterns that acquire finer spatial frequencies than A^C . For example, if measurements in A^C have a macro pixel size of 8×8 then B would contain random measurements with macro pixel size of 4×4 . The measurement matrix with the additional refined measurements can be written as

$$A^R = \begin{pmatrix} A^C \\ B \end{pmatrix}. \quad (6)$$

The refined measurements are given by $y^R = A^R x$. Again the refined solution x^R satisfies the optimality conditions

$$(A^R)^*(A^R x^R - y^R) + \alpha p^R = 0, \quad p^R \in \partial J(x^R), \quad (7)$$

where p^R is a subgradient of J at x^R and $\partial J(x^R)$ is the subdifferential of J at x^R . To be explicit the refined solution x^R contains spatial frequencies unique to the acquisition through B along with those acquired previously through A^C . To estimate the change in information of the RoI when acquired at different resolutions we would like to derive an *a-posteriori* error estimate between the coarse solution, x^C , and the refined solution, x^R , without the evaluation of the latter. This would quantify the information change across scales for the RoI. We follow the procedure for calculating error estimates in [17], [18]. From (6) we get,

$$(A^R)^* A^R = (A^C)^* A^C + B^* B. \quad (8)$$

Subtracting (7) from (5) we get,

$$\begin{aligned} &(A^C)^* A^C x^C - (A^C)^* y^C + \alpha(p^C - p^R) \\ &\quad - (A^R)^* A^R x^R + (A^R)^* 0. \end{aligned} \quad (9)$$

Using $y^C = A^C x$, $A^R x$, and the value of $(A^C)^* A^C$ from (8) we get,

$$(A^R)^* A^R (x^C - x^R) + \alpha(p^C - p^R) = B^* B (x^C - x). \quad (10)$$

Taking a scalar product of the above with $(x^C - x^R)$ yields

$$\begin{aligned} \|A^R(x^C - x^R)\|_2^2 + \alpha D_J(x^C, x^R) \\ = \langle B(x^C - x), B(x^C - x^R) \rangle, \end{aligned} \quad (11)$$

where $D_J = \langle p^C - p^R, x^C - x^R \rangle$ is the symmetric Bregman distance [17]. Applying Young's Inequality, $\langle u, v \rangle \leq \frac{1}{2}\|u\|_2^2 + \frac{1}{2}\|v\|_2^2$, to the right hand side of (11) yields

$$\begin{aligned} \|A^R(x^C - x^R)\|_2^2 + \alpha D_J(x^C, x^R) \\ \leq \frac{1}{2}\|B(x^C - x)\|_2^2 + \frac{1}{2}\|B(x^C - x^R)\|_2^2. \end{aligned} \quad (12)$$

Using the fact that, $\|Bv\|_2^2 \leq \|Bv\|_2^2 + \|A^C v\|_2^2 = \|A^R v\|_2^2$ for any v , (12) becomes

$$\frac{1}{2}\|A^R(x^C - x^R)\|_2^2 + \alpha D_J(x^C, x^R) \leq \frac{1}{2}\|y^B - Bx^C\|_2^2, \quad (13)$$

where $y^B = Bx$ are the novel measurements. The second term on the left-hand side, $D_J(x^C, x^R)$, is the error between the two solutions at coarse and fine resolutions with respect to the functional J . For $J = \|\cdot\|_1$, the symmetric Bregman distance, $D_J(x^C, x^R) = 2\sum_{\text{sign}(x_i^C) \neq \text{sign}(x_i^R)} |x_i^R - x_i^C|$, measures the deviation between entries of the two solutions that differ in their signs [17], [18]. For a wavelet norm, the symmetric Bregman distance will measure the deviation in the sparsity pattern of the wavelet representations of the two solutions [19]. If the entries of the wavelet representations have the same sign, $D_J(x^C, x^R) = 0$ and it is non-negative if the entries differ in their signs. Further, the first term on the left-hand side of (13) measures how well we estimate the magnitudes of the entries of each of the two solutions, the residual between the two solutions. The behaviour of the Bregman distance with respect to the parameter α and its asymptotic are well understood and not further discussed here, see [17]. The right-hand side of (13) is an a-posteriori estimator for the change between x^C and x^R given only the coarse solution x^C . It is an a-posteriori error estimator in the *computational* sense, it can be computed without solving a fine (the refined) scale problem. It is not an a-posteriori in the *measurement* sense, since we need to collect at least some fine scale measurements (refined measurements). Taking into consideration that B is random, we can take expectation on both sides of (13)

$$\begin{aligned} \frac{1}{2}\mathbb{E}[\|A^R(x^C - x^R)\|_2^2] + \alpha\mathbb{E}[D_J(x^C, x^R)] \\ \leq \frac{1}{2}\mathbb{E}[\|y^B - Bx^C\|_2^2]. \end{aligned} \quad (14)$$

This shows that in expectation the maximum error between x^C and x^R occurs when the expected deviation between y^B and Bx^C is maximal. Thus, the Refinement Indicator (RI) can be defined as $\|y^B - Bx^C\|_2^2$. The RI can be calculated for each RoI in the scene and the RoI with the largest value of RI can be refined. Explicitly, to refine an RoI is to calculate the

higher spatial resolution solution from the measurements y^R obtained through the measurement matrix A^R . Intuitively, for a flat region, x^C will be a good approximation of the underlying region, therefore, the novel measurements, y^B , and the *simulated* measurements obtained through B by assuming x^C to be the groundtruth; i.e. Bx^C , will be similar and the value of RI will be small. On the other hand, if the underlying region is textured, x^C will be a poor approximation to the region and the novel measurements and *simulated* measurements will differ leading to a larger value of RI. Thus, the RI can be used to prioritise RoIs for acquisition.

B. Adaptive Refinement of Measurements

The RI developed in the previous section is used to develop an acquisition algorithm for measurement-constrained systems. The three step sampling procedure discussed in [7] and further developed in Section III is used as the basic framework for acquisition. The inputs to the algorithm are the initial macro pixel size of the low resolution acquisition in the first step and a fixed number of measurements, called the measurement budget. After the ‘‘Low resolution acquisition and reconstruction’’ and ‘‘RoI detection and segmentation’’ steps we obtain the acquisition masks for each RoI. Coarse measurements for each RoI are performed and the coarse solution x^C is calculated. These coarse measurements differ from the ones acquired in the ‘‘Low resolution acquisition and reconstruction’’ step because the low resolution measurements are multiplexed measurements from the entire scene and recovering individual RoIs from multiplexed measurements is not possible. Additional random measurements y^B refining the previous coarse measurements are acquired for each RoI. The RI is calculated for all the RoIs, the RoI with the largest RI is chosen and the refined solution, x^R , is calculated for that RoI and it is set as the new coarse level for the chosen RoI. New refinement measurements are acquired at a higher spatial resolution for the current refined RoI and the new RI is calculated from the new refined measurements. RI values for the other RoIs are brought forward and again the RoI with the largest value of RI is chosen for refinement. This process continues until all the RoIs are in the native resolution of the DMD or until the measurement budget is exhausted. This results in prioritisation of RoIs in terms of the information content change across spatial resolutions. The limited measurements are spent on RoIs with greater change in information content. If measurements are not available for refining a particular RoI then it is removed from the prioritisation list and the RIs of the remaining RoIs are compared for refinement. The complete algorithm is shown in Fig. 6 and is named RoI Prioritised Sampling (RPS).

C. Discussion

One can consider a naive approach in which one acquires RoIs at the coarse resolution, performs coarse reconstructions and acquires further refined measurements for all the RoIs. One can reconstruct all the RoIs at this refined resolution and then select the RoI that produced maximum change in information content between the coarse and the refined image of the RoI

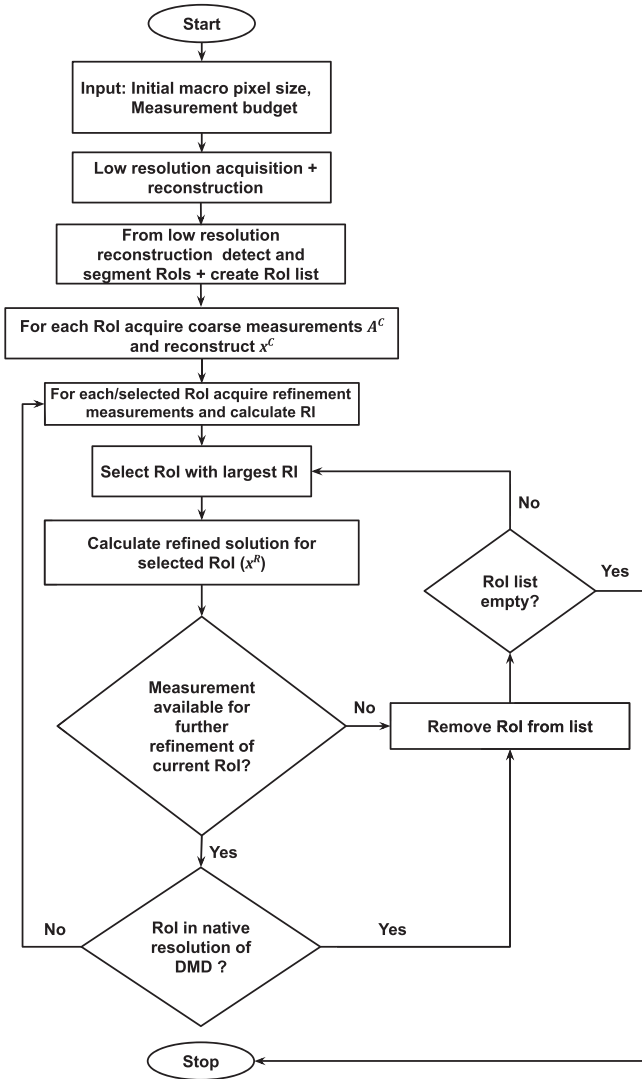


Fig. 6. RPS Algorithm for ROI prioritisation using RI in measurement-constrained autonomous exploration systems.

for further refinement, expecting the trend of increase in information content over scales to continue. However, the decision on which ROI should be refined further can only be made *after the reconstruction of all ROIs at the refined resolution*. The extra reconstructions may be expensive for a constrained system. This problem is alleviated with the RPS, as RI allows us to make the same decision *without* calculation of the refined solution for *all* ROIs. In limited measurement settings, like remote exploration, acquired refined measurements can be sent to a ground-station where resources would be available for reconstructions. Thus, we focus specifically on a measurement-constrained setting. Also, the RPS is concerned with utilising the limited measurement budget optimally and not specifically with reducing the number of measurements required to acquire the scene/ROI.

V. RESULTS

To demonstrate that the RI does in fact quantify information content change at different spatial scales, three different regions

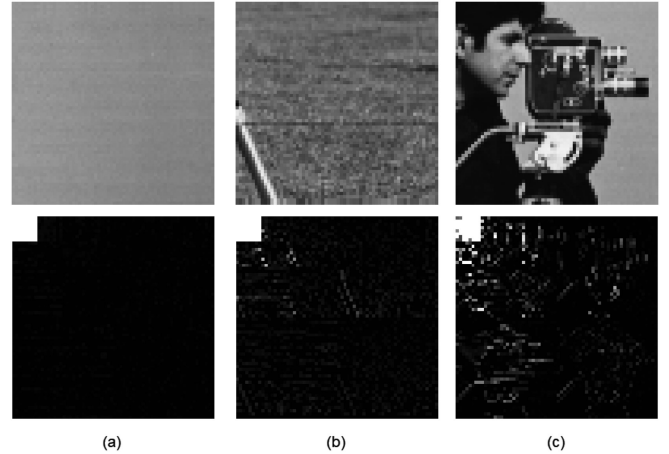


Fig. 7. *Top Panel*: Three different ROIs from the cameraman image with different amount of textures. *Bottom Panel*: The three-level Haar wavelet decomposition of the corresponding ROIs. (a) ROI with minimal textures (flat region), (b) ROI with some textures, and (c) ROI with rich textures.

TABLE I
RI VALUES FOR ROIS FROM FIG. 7 CALCULATED AT SPATIAL RESOLUTIONS OF 8×8 , 4×4 , AND 2×2 MACRO PIXELS. RANDOM-MACRO-PIXEL MEASUREMENTS WITH 0/1 ENTRIES WERE USED FOR EACH SPATIAL RESOLUTION. THE NUMBER OF REFINED MEASUREMENTS AT EACH MACRO PIXEL RESOLUTION FOR EACH ROI IS 409

Spatial Resolution	RI for ROI Fig. 7(a)	RI for ROI Fig. 7(b)	RI for ROI Fig. 7(c)
8×8	9.74	97.91	218.11
4×4	5.79	61.58	110.37
2×2	5.10	39.34	46.99

of the cameraman image were selected with varying amount of textures. The ROIs and their three-level Haar wavelet decompositions are shown in Fig. 7. Random-macro-pixel measurements with 0/1 entries are performed on each ROI starting from a macro pixel size of 8×8 and the 8×8 reconstruction is calculated. Thereafter, a refinement matrix B is generated by using random-macro-pixel sampling with a macro pixel size of 4×4 and novel measurements are generated for each ROI. *Simulated* measurements are also generated using the reconstruction of the 8×8 macro pixel acquisition and the refinement matrix B . The two measurements are subtracted and the l_2 norm of the error vector is calculated to form the RI. The process is repeated for macro pixel size of 2×2 . The RI values for the ROIs are recorded in Table I. The RI for spatial resolution 8×8 denotes the information change in a ROI when measurements are refined from a resolution of 8×8 to 4×4 . The last row in the table defines the RI value of refining a ROI from 2×2 to 1×1 , the native resolution of the DMD, and there is no RI for spatial resolution of 1×1 macro pixel. The largest change in the value of the RI is consistently obtained for the ROI that has the largest number of high frequency components, i.e., for ROI (c) in Fig. 7. The smallest change in RI value is observed for ROI (a) that has the smallest number of high frequency components. Thus, the RI estimates the change in information content across spatial scales and is a relevant metric that can be used to prioritise ROIs for acquisition in measurement-constrained systems.

Further, the RPS algorithm developed in the previous section was tested on images derived from different sensors. The experiments are designed to reflect remote exploration scenarios from Earth observation applications and extra-terrestrial exploration. After a short overview of the datasets the results of applying RPS on the datasets will be discussed and the efficacy of the proposed method in prioritising RoIs in measurement-constrained systems will be verified. The algorithm was implemented in MATLAB 2018a and simulations were performed on a Windows system with 16 GB RAM and an Intel(R) i7-6700 CPU @ 3.4 GHz. Code for reproducing the results can be found at <https://github.com/protim1191/RoI-Prioritised-Sampling.git>.

A. Datasets

1) *Cuprite*: The Cuprite geological dataset is a snapshot of the cuprite ores in the state of Nevada, USA. The hyperspectral datacube was acquired by the AVIRIS [35] sensor that collects data in the wavelength range of 400 to 2500 nm with a nominal channel bandwidth of 10 nm. The ground sampling area is 20 m² and the radiometric resolution is 16 bits. The datacube is divided into two, the visible and the infrared. The visible section consists of wavelengths 400 to 800 nm and is called the “Cuprite Vis” dataset; the infrared dataset consists of wavelengths from 900 to 2500 nm and is called the “Cuprite IR” dataset.

2) *Gulf of Mexico*: The Gulf of Mexico dataset is a part of baseline datasets provided by SpecTIR [36]. The sensor acquires data in the wavelength range of 395 to 2450 nm with a nominal channel bandwidth of 5 nm. The ground sampling distance is 2 m. As before the datacube is divided into visible and infrared regions. The visible region consists of wavelengths from 400 to 800 nm and is called the “Gulf Vis” dataset. The infrared region consists of wavelengths from 900 to 2400 nm and is called “Gulf IR” dataset. Only the Gulf Vis dataset is used for experiments as the spatial variations in the two datasets are similar.

3) *University of Pavia*: This scene was acquired by the RO-SIS sensor [37] during a flight campaign over Pavia in northern Italy. The detector is sensitive to wavelengths ranging from 430 to 860 nm with a nominal bandwidth of 5 nm. Certain pixels did not contain any information and were discarded. The entire dataset is considered to be a part of the visible wavelengths and is called “Pavia” in this work.

4) *MSL ChemCam Images*: Two images from the Mars Science Laboratory (MSL) [38] on the Curiosity rover *0133_crc_ccam01133_rowatt.png* and *0121_crc_ccam01121_stanbridge.png* have been used to test the RPS algorithm. The test images are called “Mars1” and “Mars2” respectively. These are shown in Fig. 9.

B. Preprocessing of Datasets and Algorithm Parameters

For all the datasets a 256 × 256 section was used. After removing corrupted bands, pixel values in each spectral band were converted to be between 0 and 1 by using min-max normalisation. A single channel image was formed from the normalised spectral image by averaging across the spectral dimension of each pixel. The SPC architecture described in Section III-A was used to simulate the acquisition process. The measurement

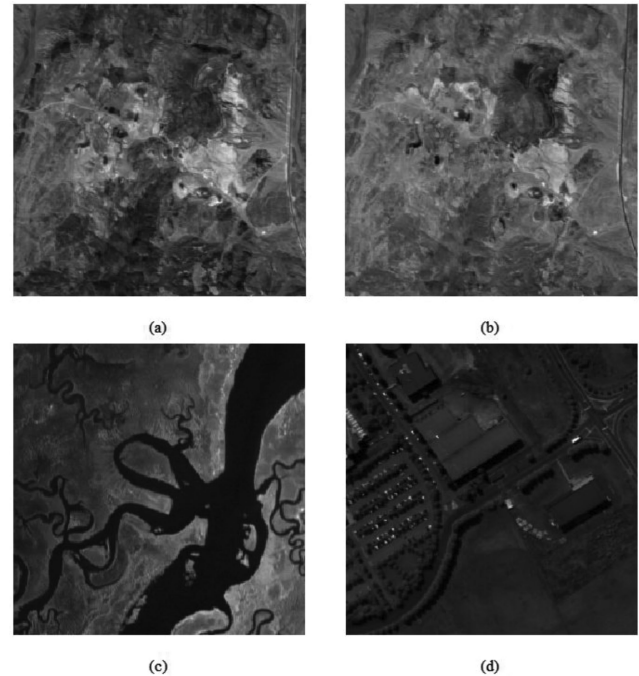


Fig. 8. Remote sensing datasets used for evaluation of RPS. (a) Cuprite Vis dataset, (b) Cuprite IR dataset, (c) Gulf Vis dataset, and (d) Pavia dataset.

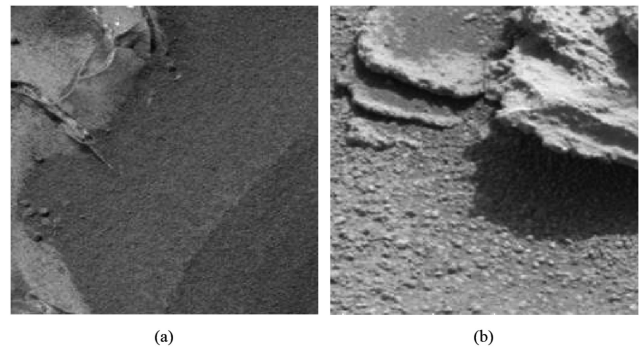


Fig. 9. Datasets from the Mars Science Laboratory on the Curiosity rover used for evaluation of RPS. (a) Mars1 and (b) Mars2.

budget was different for each dataset and is specified along with their respective results. The number of low resolution measurements in the first step was fixed at 1000 and the macro pixel size used was 8×8 . The number of regions selected during the “RoI detection and segmentation” procedure was fixed at 10. Regions detected within 1 pixel radius of each other were merged. Thus, the final number of RoIs available for refinement is different for each dataset. The acquisition and recovery of each RoI is independent of the other RoIs, therefore, the overlap regions common to two RoIs is sampled twice at the same resolution. However, due to budget considerations if only one of the RoIs can be acquired, then the overlap region will be at higher resolution for the unacquired RoI as well. As the algorithm keeps track of the resolution of each RoI, further visualization or analysis of the overlap region can be restricted to the higher resolution RoI. The number of measurements used to calculate the RI

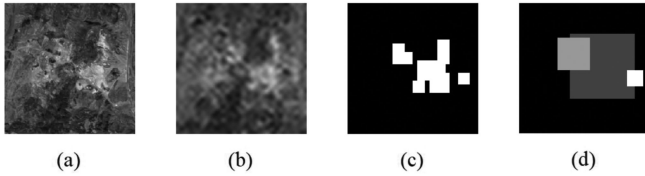


Fig. 10. Low resolution reconstruction and RoI detection and segmentation process for Cuprite Vis dataset. (a) Original, (b) low resolution reconstruction from 1000 measurements, (c) detected and segmented RoIs, and (d) merged RoIs. The three gray-levels represent three different RoIs for which the RI will be calculated.

was fixed at 10% of the total number of pixels in the RoI for random-macro-pixel sampling and at 20% for multi-scale multi-level sampling. The RoI with the largest value of RI is chosen for refinement at each step.

Upon experimentation algorithm (1) was found to be better at the initial low resolution reconstruction of the entire scene with 2D-DCT as the sparsity basis. For subsequent reconstruction of individual RoIs using RPS, algorithm (2) was used with Daubechies-8 wavelets [39] as the sparsity basis. The algorithms were implemented in the TFOCS [32] toolbox. The maximum number of iterations was fixed at 20000, the algorithm used was ‘N07,’ continuation was set to ‘True’ and the number of continuation iterations were set to 3, the regularisation parameters β_1 and β_2 from (2) were set to 1 and 0.4 respectively for random-macro-pixel sampling and $\beta_1 = 1$ and $\beta_2 = 0.6$ for Walsh measurements. The measurement and sparsity operators were implemented in the SPOT toolbox [40]. The results of “Low resolution acquisition and reconstruction” and “RoI detection and segmentation” process for Cuprite Vis dataset is shown in Fig. 10. Results after segmentation are discussed next.

C. Results on Remote Sensing Datasets

This section discusses the results for the remote sensing datasets, i.e., Cuprite Vis, Cuprite IR, Gulf Vis, and Pavia, shown in Fig. 8. For the following results multi-scale random-macro-pixel sampling with 0/1 entries was used. The results of applying the RPS algorithm on the Cuprite Vis dataset is shown in Fig. 11. The first image from the left in the top row is the low resolution reconstruction. The second image from the left in the top row shows the coarse reconstruction of all the detected RoIs. Continuing from left to right and top to bottom, the prioritisation of the RoIs and their refined reconstructions with the progress of the RPS is shown. Table II records the change in the RI values and the choice of RoI made at each iteration of the algorithm along with the number of available measurements at each iteration that decides whether further refinement is possible for a RoI. For the Cuprite Vis dataset a total of 9600 measurements was used for acquisition, which is 14.6% of the total number of pixels in the scene. The number of low resolution measurements is 1000 and these are used to acquire the entire scene at 8×8 macro pixel resolution. Three RoIs of size 128×128 , 64×64 , and 32×32 are selected after the “RoI detection and segmentation” process (Fig. 10). A total of 2149 ($1638 + 409 + 102$) coarse measurements were used.

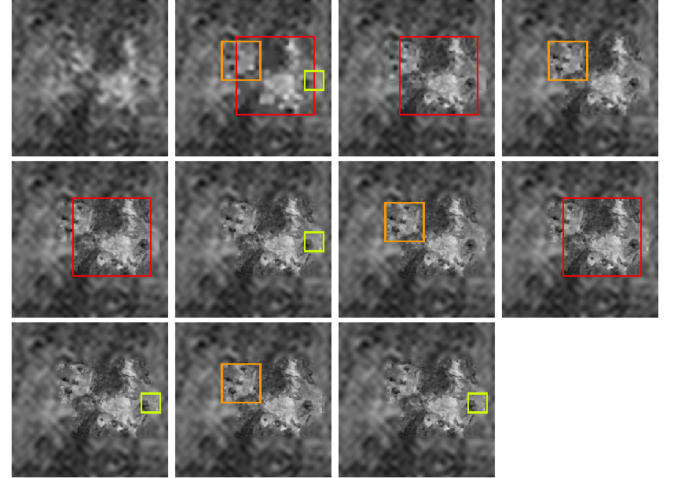


Fig. 11. Prioritisation of RoIs with RI for Cuprite Vis dataset. The total number of measurements is 9600 (14.6% of the total number of pixels).

TABLE II
EVOLUTION OF RI FOR CUPRITE VIS DATASET. THE HIGHLIGHTED ROI INDICATES THE ONE SELECTED AT EACH ITERATION. THE REFINED MACRO PIXEL RESOLUTION OF THE SELECTED ROI IS TABULATED. THE TOTAL NUMBER OF MEASUREMENTS IS 9600

Iteration No.	RI			Refined Macro Pixel Resolution	Available Measurements
	RoI 1 (Red) (128×128)	RoI 2 (Orange) (64×64)	RoI 3 (Yellow) (32×32)		
1	349.30	134.64	14.82	4×4	4302
2	56.98	134.64	14.82	4×4	2664
3	56.98	14.64	14.82	2×2	2255
4	11.73	14.64	14.82	4×4	617
5	11.73	14.64	6.57	2×2	515
6	11.73	3.36	3.36	1×1	106
7	-	3.36	6.57	2×2	106
8	-	3.36	0.80	1×1	4
9	-	-	0.8	1×1	4
10	-	-	-	-	4

Another 2149 measurements were used for the initial RI values for the three RoIs. These values indicate the change in information content for each RoI when refined to a macro pixel resolution of 4×4 from their current 8×8 macro pixel resolution. The 128×128 sized RoI, marked in red in Fig. 11, is selected for refinement as it has the largest RI value. Upon calculation of the refined solution of the selected RoI the number of available measurements is checked to decide whether enough measurements is available for calculation of the new RI for the selected RoI, which would estimate the change in information one could expect when the RoI is refined to a macro pixel resolution of 2×2 from the current 4×4 . As there are 4302 available measurements at iteration 1 and 1638 measurements are required for the selected RoI, refined measurements are performed on the selected RoI and the new RI is calculated. At the second iteration the same sequence of steps is followed, the RoI with the largest value of RI is selected for refinement and the new RI is calculated subject to the availability of measurements. The last RI calculation is done for estimating the change in

TABLE III
ERROR METRICS FOR THE ROIS FOR CUPRITE VIS DATASET AT VARIOUS
MACRO PIXEL RESOLUTION

Resolution of RoI	RoI 1 (Red) (128 × 128)		RoI 2 (Orange) (64 × 64)		RoI 3 (Yellow) (32 × 32)	
	NMSE	SSIM	NMSE	SSIM	NMSE	SSIM
8 × 8	0.032	0.40	0.041	0.36	0.024	0.38
4 × 4	0.015	0.64	0.016	0.64	0.015	0.56
2 × 2	0.008	0.79	0.009	0.79	0.006	0.78
1 × 1	0.007	0.83	0.007	0.83	0.005	0.82

information when refining from 2×2 to 1×1 macro pixel resolution. If the number of available measurements allowed this acquisition then the 1×1 solution is always calculated, which may also be performed at the ground-station as the measurements for this reconstruction step have already been acquired for the calculation of the RI. Once a RoI is resolved to 1×1 macro pixel resolution it is replaced with a '-' in the table. This can be seen in iteration 6 and 7 in Table II where at iteration 6 the red RoI is selected for refinement to a macro pixel resolution of 1×1 and replaced with a '-' in iteration 7 but no measurements are utilised as no further refinement is possible for this RoI. As a result, the same number of measurements is available at iteration 7 for RI calculations as at iteration 6.

The refinement measurements used to calculate the RI are also used for calculating the refined solution. Table III records the quality of reconstruction of the selected RoIs at different macro pixel resolutions. The quality of recovery for each RoI improves with acquisition of higher spatial frequency components. Two error metrics are used, namely the Normalised Mean Squared Error (NMSE) and the Structural SIMilarity Index (SSIM) [41]. NMSE between a vector x and its estimate \hat{x} is defined as $\|\hat{x} - x\|_2^2 / \|x\|_2^2$. SSIM is a full referential perceptual visual quality metric that considers local luminance, contrast, and structural variance to calculate similarity between the estimate and the reference image. The visual improvements in the reconstructions can be seen in Fig. II in the supplementary material.

Results of applying RPS on the Cuprite IR dataset are shown in Fig. 12. As shown in Table IV, the online distribution of measurements is more evident here. After coarse reconstruction the RI values for each RoI is calculated. The RoI with the largest RI value, marked in red, is selected for refinement. The available number of measurements allows to calculate the RI for a macro pixel resolution of 2×2 . In the second iteration, the RoI of size 64×64 , marked in orange in Fig. 12, has the largest value of RI and is refined to a macro pixel resolution of 4×4 . In the third iteration, the red RoI has the largest RI value and is refined to a macro pixel resolution of 2×2 . However, there is not enough measurements to calculate the RI for a refinement to a macro pixel resolution of 1×1 . Therefore, the red RoI is removed from the list of RoIs due to unavailability of measurements and replaced with a '-'. The same number of measurements is available at iteration 4 as at iteration 3 because no refinement measurements were used in iteration 3. Thereafter, the process of selecting the RoI with the largest value of RI, calculation of the refined solution, and calculation of the RI subject to

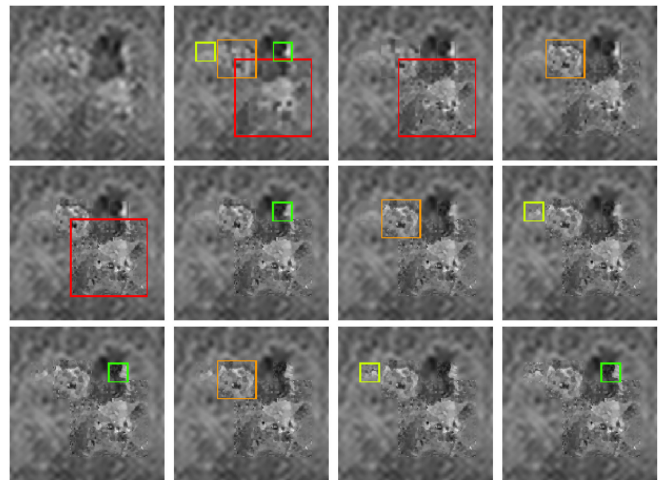


Fig. 12. Prioritisation of RoIs with RI for Cuprite IR dataset. The total number of measurements is 8060 (12.3% of the total number of pixels).

TABLE IV
EVOLUTION OF RI FOR CUPRITE IR DATASET IN FIG. 12. THE HIGHLIGHTED
ROI INDICATES THE ONE SELECTED AT EACH ITERATION. THE REFINED MACRO
PIXEL RESOLUTION OF THE SELECTED ROI IS TABULATED. THE TOTAL
NUMBER OF MEASUREMENTS IS 8060

Iteration No.	RI				Refined Macro Pixel Reso- lution	Available Mea- sure- ments
	RoI 1 (Red) (128×128)	RoI 2 (Orange) (64 × 64)	RoI 3 (Green) (32 × 32)	RoI 4 (Yellow) (32 × 32)		
1	272.21	139.14	8.39	29.55	4 × 4	2762
2	41.43	139.14	8.39	29.55	4 × 4	1124
3	41.43	14.04	8.39	29.55	2 × 2	715
4	-	14.04	8.39	29.55	4 × 4	715
5	-	14.04	8.39	4.13	2 × 2	613
6	-	3.44	8.39	4.13	4 × 4	204
7	-	3.44	1.80	4.13	2 × 2	102
8	-	3.44	1.80	1.10	1 × 1	0
9	-	-	1.80	1.10	2 × 2	0
10	-	-	-	1.10	1 × 1	0
11	-	-	-	-	-	0

available measurements continues till iteration 7. At iteration 7 the yellow RoI uses the last measurements to calculate the RI for a refinement to a macro pixel resolution of 1×1 . At iteration 8 the orange RoI is resolved to 1×1 macro pixel resolution as the required measurements were already acquired at iteration 5. At iteration 9 the green RoI is resolved to a macro pixel resolution of 2×2 but further refinement is not possible as the measurement budget has been exhausted. Therefore, the final macro pixel resolution of the green RoI is 2×2 . Finally, the yellow RoI is resolved to the native resolution of the DMD by using the refinement measurements collected at iteration 7. This shows the dynamic distribution of the measurements across the RoIs by the RPS algorithm. For acquisition of all the RoIs at the native resolution of the DMD, one would need 10000 measurements. Thus, the RPS algorithm selects the most informative RoI at each iteration and better utilises the limited measurement budget (8060 measurements).

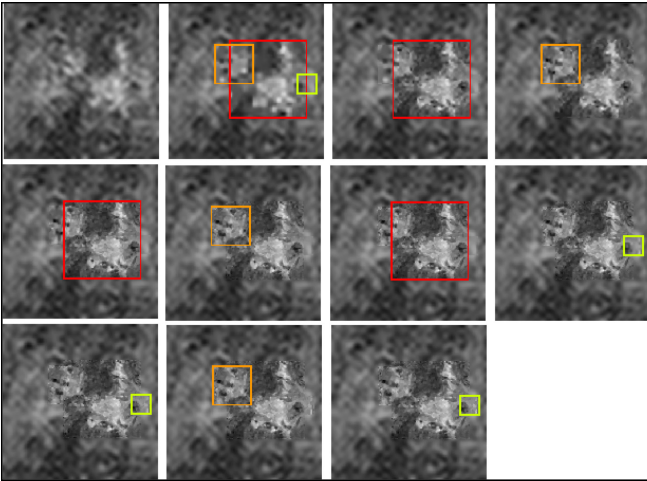


Fig. 13. Prioritisation of RoIs with RI for Cuprite Vis dataset with random-macro-pixel Rademacher measurements. The total number of measurements is 9188 (14% of the total number of pixels). The number of physical measurements is 18376 (28% of the total number of pixels).

TABLE V
ERROR METRICS FOR THE ROIS FOR CUPRITE VIS DATASET WITH
RANDOM-MACRO-PIXEL RADEMACHER MEASUREMENTS (FIG. 13).
RESOLUTION IS IN MACRO PIXEL SIZE

Resolution of RoI	RoI 1 (Red) (128 × 128)		RoI 2 (Orange) (64 × 64)		RoI 3 (Yellow) (32 × 32)	
	NMSE	SSIM	NMSE	SSIM	NMSE	SSIM
8 × 8	0.032	0.39	0.043	0.35	0.022	0.4
4 × 4	0.015	0.63	0.017	0.63	0.015	0.54
2 × 2	0.008	0.78	0.009	0.79	0.007	0.77
1 × 1	0.006	0.83	0.007	0.83	0.005	0.83

Due to space limitations the application of RPS on the Gulf Vis dataset is shown in Fig. III and Table A in the supplementary material. The results for Pavia are provided in Fig. IV and Table B in the supplementary material. In the case of Pavia we can see that a smaller RoI (RoI 1 and RoI 4) with more information change across the spatial scales is prioritised over a larger sized RoI (RoI 3).

D. Experiments With Different Measurement Matrices

The RPS is agnostic to the nature of the measurement matrix. The only requirement is that the measurement ensemble can be deployed in multi-scale manner. Fig. 13 shows the result of application of the RPS on the Cuprite Vis dataset with Rademacher (± 1 with equal probability) [42] random-macro-pixel measurements. The RoIs and their order of prioritisation changes, though the overall result is similar to 0/1 random-macro-pixel sampling. The corresponding error metrics are shown in Table V. The first step of low resolution acquisition for detecting RoIs is still performed with 0/1 random-macro-pixel sampling. Implementation of Rademacher random-macro-pixel measurements requires two acquisition cycles on the DMD due to the ± 1 nature of the Rademacher random variable. The total number of physical measurements required for Rademacher measurements is twice the number of realisations of the measurement vectors.

TABLE VI
NUMBER OF MEASUREMENTS PER SPATIAL FREQUENCY BAND IN WALSH
SAMPLING MASKS

Mask Size	Number of Required Measurements		
	Low Frequency	Medium Frequency	High Frequency
128 × 128	2025	689	562
64 × 64	676	79	64
32 × 32	121	67	16

Let the number of measurements required for calculation of RI at each spatial scale for the i^{th} RoI be $p_i\%$ of the total number of pixels in the RoI. Then, for M RoIs with R resolution levels, the total number of 0/1 random-macro-pixel measurements will be $CM + R \sum_{i=1}^M \lfloor \frac{p_i}{100} n_i \rfloor$, where CM is the number of low resolution measurements common for all RoIs and n_i denotes the total number of pixels in each RoI. For Rademacher random-macro-pixel measurements, the total number of physical measurements, i.e., the number of acquisition cycles performed on the DMD, will be equal to $CM + 2R \sum_{i=1}^M \lfloor \frac{p_i}{100} n_i \rfloor$.

One can also use Walsh transforms in a multi-scale manner. The Walsh sampling masks from Fig. 5 are used. The number of Walsh measurements in each spatial frequency band is provided in Table VI. For each RoI size the total number of Walsh measurements is fixed beforehand while determining the sampling map. The number of measurements required for RPS with Walsh ensemble is equal to $CM + \sum_{i=1}^M \lfloor \frac{p_i}{100} n_i \rfloor$. The total physical measurements cycles on the DMD is twice the second term, as the Walsh ensemble consists of ± 1 measurements, plus the low resolution measurements for the entire scene. The low resolution acquisition is performed with 0/1 random-macro-pixel measurements. For RoI reconstruction only the Walsh measurements are used. For calculating the initial RI values, the low frequency Walsh measurements are used as the refined measurements and the low resolution regions corresponding to the RoIs from the low resolution reconstruction are used as the coarse reconstruction for the *simulated* measurements. For further refinement the Walsh reconstructions at different spatial scales are used as the coarse reconstructions for the *simulated* measurements. The results with multi-level Walsh measurements are shown in Fig. 14 and in Table VII. The reconstruction metrics are recorded in Table VIII. The advantage of using the concept of multi-level sampling with the concept of RI is evident from comparing the results in Table V and Table VIII. Multi-level Walsh measurements with RI achieve similar performance as random-macro-pixel Rademacher measurements with a smaller measurement budget.

E. Results on MSL Images

This section reports the results of using RPS to reconstruct the Mars1 and Mars2 images from Walsh measurements. The prioritisation of RoIs for Mars1 image is shown in Fig. 15. The evolution of RI values and RoI prioritisation are recorded in Table IX. A limited budget experiment is also conducted with the Mars1 image where the number of measurements is not enough to resolve each RoI to the native resolution of the DMD. Fig. 16 and Table X show the prioritisation of the RoIs

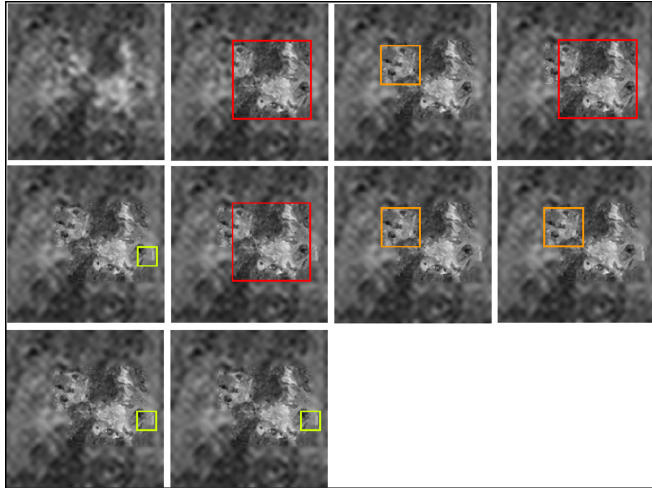


Fig. 14. Prioritisation of RoIs with RI for Cuprite Vis dataset with Walsh measurements exploiting multi-level sampling. The total number of measurements is 5300 (8.1% of the total number of pixels). The number of physical measurements is 9600 (14.6% of the total number of pixels).

TABLE VII
EVOLUTION OF RI FOR CUPRITE VIS DATASET WITH WALSH MEASUREMENTS USING MULTI-LEVEL SAMPLING (FIG. 14). THE HIGHLIGHTED ROI INDICATES THE ONE SELECTED AT EACH ITERATION

Iteration No.	RI			Available Measurements
	RoI 1 (Red) (128 × 128)	RoI 2 (Orange) (64 × 64)	RoI 3 (Yellow) (32 × 32)	
1	520.56	110.12	24.83	1478
2	81.84	110.12	24.83	789
3	81.84	18.15	24.83	710
4	22.44	18.15	24.83	148
5	22.44	18.15	5.50	81
6	-	18.15	5.50	81
7	-	7.30	5.50	17
8	-	-	5.50	17
9	-	-	3.54	1
10	-	-	-	1

and the evolution of the RI values respectively. The orange RoI is resolved using only the low frequency Walsh measurements denoted by the green region in Fig. 5(c) and the corresponding number of measurements is given in Table VI. As there is not enough measurements for further refinement of the orange RoI, it is not resolved using the higher frequency Walsh measurements and is removed from the list. Only the red RoI is resolved to the native resolution of the DMD. Thus, the RI is able to prioritise RoIs on the basis of the information change across spatial scales of the RoIs from Mars rover camera images as well. Error metrics for both measurement budget scenarios of the Mars1 image are provided in Table C and Table D in the supplementary material. The result of applying the RPS on the Mars2 image is shown in Fig. 17. The evolution of the RI and RoI selection is shown in Table E in the supplementary material. The error metrics are recorded in Table F in the supplementary material.

TABLE VIII
ERROR METRICS FOR THE ROIS FOR CUPRITE VIS DATASET WITH WALSH MEASUREMENTS (FIG. 14) USING MULTI-LEVEL SAMPLING IN THE RPS ALGORITHM

	RoI 1 (Red) (128 × 128)		RoI 2 (Orange) (64 × 64)		RoI 3 (Yellow) (32 × 32)	
	NMSE	SSIM	NMSE	SSIM	NMSE	SSIM
Coarse Step	0.031	0.42	0.031	0.45	0.022	0.41
Refinement Step 1	0.008	0.79	0.007	0.82	0.009	0.71
Refinement Step 2	0.007	0.82	0.007	0.83	0.007	0.76
Refinement Step 3	0.006	0.83	0.007	0.84	0.006	0.77

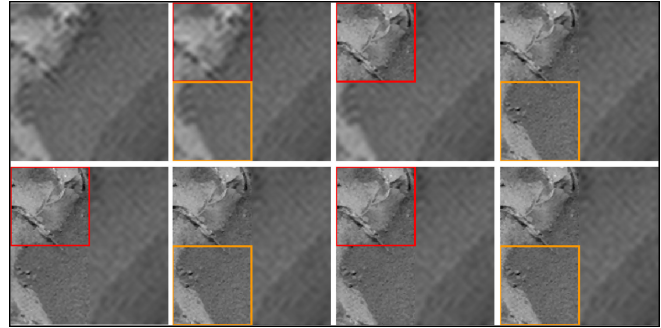


Fig. 15. Prioritisation of RoIs with RI for Mars1 image. The total number of measurements is 7552 (11.5% of the total number of pixels).

TABLE IX
EVOLUTION OF RI FOR MARS1 IMAGE WITH WALSH MEASUREMENTS USING MULTI-LEVEL SAMPLING (FIG. 15). THE HIGHLIGHTED ROI INDICATES THE ONE SELECTED AT EACH ITERATION. THE TOTAL NUMBER OF MEASUREMENTS IS 7552

Iteration No.	RI		Available Measurements
	RoI 1 (Red) (128 × 128)	RoI 2 (Orange) (128 × 128)	
1	293.72	147.39	2502
2	64.84	147.39	1813
3	64.84	49.33	1124
4	14.37	49.33	562
5	14.37	13.75	562
6	-	13.75	0
7	-	-	0

TABLE X
EVOLUTION OF RI FOR MARS1 IMAGE WITH WALSH MEASUREMENTS USING MULTI-LEVEL SAMPLING (FIG. 16) WITH LIMITED MEASUREMENT BUDGET. THE HIGHLIGHTED ROI INDICATES THE ONE SELECTED AT EACH ITERATION. THE TOTAL NUMBER OF MEASUREMENTS IS 6301

Iteration No.	RI		Available Measurements
	RoI 1 (Red) (128 × 128)	RoI 2 (Orange) (128 × 128)	
1	293.72	147.39	1251
2	64.84	147.39	562
3	64.84	-	562
4	14.37	-	562
5	-	-	0

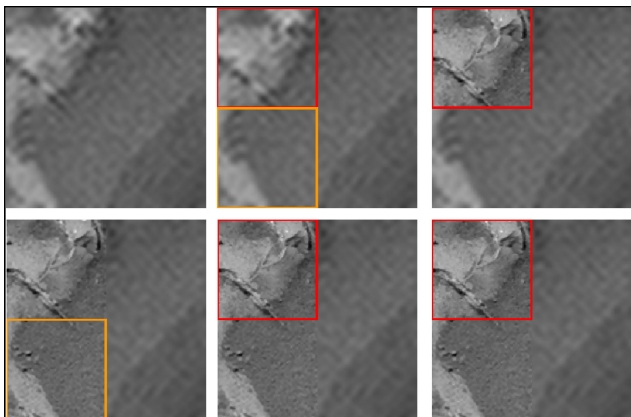


Fig. 16. Prioritisation of RoIs with RI for Mars1 image with limited measurement budget. The total number of measurements is 6301 (9.6% of the total number of pixels).

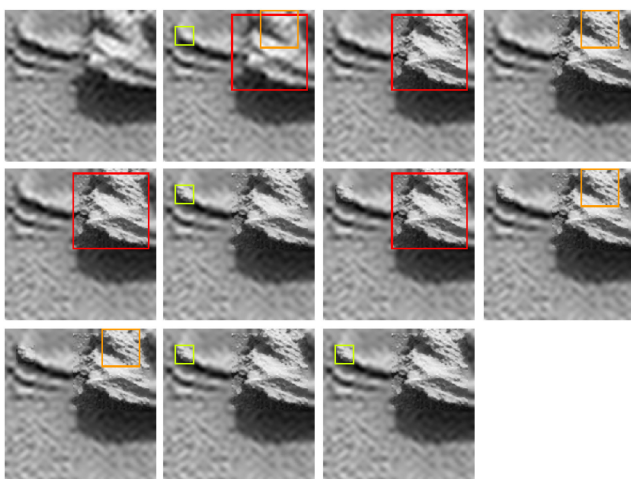


Fig. 17. Prioritisation of RoIs with RI for Mars2 image. The total number of measurements is 5299 (8.1% of the total number of pixels).

F. Discussion: Noisy Measurements and Processing Time

In case of noisy observations, the RPS is still applicable. The error estimate can be used in the same way, since the RI only takes into account the given data independent of their model. The reason is that we approximate regularised solutions at finer levels (that is an estimator based on the given noisy data) instead of an idealised solution. Consequently, the output of the “detection and segmentation” step may lead to different RoI selections with the noise realisations, but we expect the first ones to remain more stable. Deviations from an idealised solution are rather a model error than due to our adaptive algorithm, the issue of simultaneously treating modelling errors is beyond the scope of the paper and left to future research. Examples of RPS with noisy observations using the Cuprite Vis dataset are provided in Fig. 5, 6, 7, and 8 in the supplementary material. In case of low light conditions, if the photon count is not too low then the noise can be modelled as a Gaussian with a variable variance and the RI is applicable. Analysing scenarios dominated by Poisson

TABLE XI
ERROR METRICS FOR COMPARISON OF RPS WITH CLASSICAL AND MULTI-LEVEL COMPRESSED SENSING (CS)

		Classical CS	Multi-level CS	RPS
NMSE	RoI 1 (Red)	0.016	0.011	0.006
	RoI 2 (Orange)	0.017	0.011	0.007
	RoI 3 (Yellow)	0.016	0.011	0.006
	Background	0.029	0.012	0.028
SSIM	RoI 1 (Red)	0.61	0.71	0.83
	RoI 2 (Orange)	0.65	0.73	0.84
	RoI 3 (Yellow)	0.65	0.71	0.77
	Background	0.7	0.79	0.62

noise is beyond the scope of the article, it will be considered in future work along the lines of [43].

The required processing time depends on the underlying compressed sensing solver. This can be tailored to the required application and may be iterative or learned in nature. If the application admits structured measurement matrices, such as the Walsh or Fourier matrices, then the FFT algorithm can be exploited to accelerate the reconstruction. The complexity analysis of compressed sensing solvers can be found in [32] and the references therein. Calculation of the RI involves a vector subtraction that takes N FLOPs and a calculation of a squared l_2 norm that takes $2N - 1$ FLOPs, where N is the total number of pixels in the RoI. The total number of FLOPs required for RI calculation is $3N - 1$.

G. Comparison With Other Techniques

As discussed previously, for exploration scenarios region based acquisition is preferred over acquisition of the complete scene due to limited resources. Also a region based approach provides the opportunity to acquire the RoIs with greater precision. A comparison of the RPS with classical compressed sensing [8], [9] and multi-level compressed sensing [13] with Walsh measurements is done in Table XI. The measurement budget is 5300 for all three acquisition methods and algorithm (2) is used with Daubechies-8 wavelets as the sparsity basis.

Rademacher measurements are used for classical compressed sensing. On its own classical compressed sensing and multi-level compressed sensing recover the entire scene at the native resolution of the DMD. As expected the RoIs are best reconstructed by the RPS, as in that case the RoIs are recovered in a dedicated manner. As the background is left at a lower resolution in the RPS algorithm the recovery of the background is best with multi-level compressed sensing. Table XI shows that in an exploration scenario with limited measurement budget, region based acquisition, like the RPS, is able to distribute the measurement budget intelligently and acquire better quality RoIs, as measured by NMSE and SSIM. Fig. 9 in the supplementary material shows the reconstruction comparison from the three considered methods.

Fig. 18 and Fig. 19 show the multi-resolution compressed sensing reconstructions from MR-AMP [4], [44] and the STOne

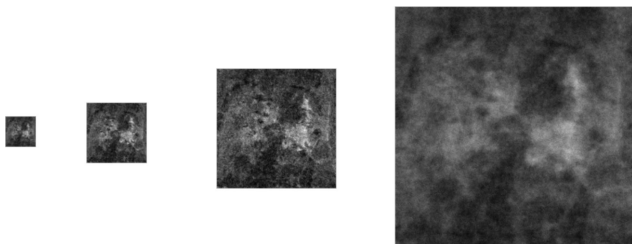


Fig. 18. MR-AMP reconstruction for Cuprite Vis dataset. From left to right down-sampling factor: 8, 4, 2, and 1. The total number of measurements is 9600 (14.6% of the total number of pixels).

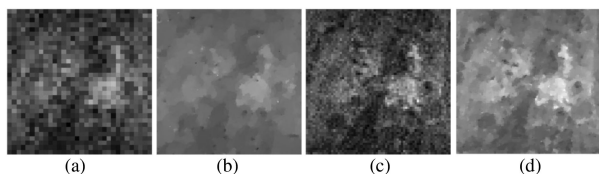


Fig. 19. STOne transformation reconstruction for Cuprite Vis dataset. (a) Low resolution reconstruction with 1024 measurements, (b) high resolution reconstruction of (a), low resolution reconstruction with 4096 measurements, and (d) high resolution reconstruction of (c).

transform [3], [45] respectively for the Cuprite Vis dataset. MR-AMP used the same number of measurements for reconstruction at different down-sampling factors as the RPS algorithm illustrated in Fig. 11. For STOne transformation the number of measurements must be a power of 2, therefore, we use 1024 and 4096 measurements to reconstruct two different low resolution images. These comparisons show that at the considered sub-sampling rates, a step-wise resolution refinement procedure, like the RPS, works better. No codes were provided by the authors of [5], therefore, no comparison was possible.

VI. CONCLUSION

This work proposes a procedure for RoI prioritisation in constrained systems. An estimator for the change in information content across spatial resolutions is developed. The Refinement Indicator (RI) evaluates the importance of a RoI for refinement to a finer resolution by calculating the error between the novel refined measurements and *simulated* high resolution measurements obtained by considering the coarse resolution reconstruction as the groundtruth. Thus, the calculation of the refined higher spatial resolution reconstruction is not required to estimate the change in information content across scales. The RI is used to prioritise RoIs for sequential acquisition through the RPS algorithm. Experiments show that the RPS algorithm is able to distribute the limited measurement budget in constrained systems in an organized manner across the different RoIs. At the considered sampling rates, the RoI based step-wise resolution refinement performs better at acquiring mission relevant RoIs than the state-of-the-art multi-resolution reconstruction techniques that recover the complete scene at a low resolution. Multi-scale multi-level compressed sensing with Walsh measurements is employed in the RPS framework and provides similar performance to random-macro-pixel sampling using Rademacher measurements with smaller measurement

budgets. The RI thus provides a method for dynamic measurement budget allocation in constrained systems, which increases the capabilities of autonomous systems in exploration scenarios.

In the future, the work may be extended to multispectral scenarios and to low photon count environments. Methods for sequential compressed sensing reconstruction [46] may be studied for this application.

REFERENCES

- [1] Y. Gao and S. Chien, "Review on space robotics: Toward top-level science through space exploration," *Sci. Robot.*, vol. 2, no. 7, Jun. 2017, Art. no. ean5074.
- [2] R. Francis *et al.*, "AEGIS autonomous targeting for the curiosity Rover's ChemCam instrument," in *Proc. IEEE Appl. Imagery Pattern Recognit. Workshop*, 2015, pp. 1–5.
- [3] T. Goldstein, L. Xu, K. F. Kelly, and R. Baraniuk, "The Stone transform: Multi-resolution image enhancement and compressive video," *IEEE Trans. Image Process.*, vol. 24, no. 12, pp. 5581–5593, Dec. 2015.
- [4] X. Wang and J. Liang, "Multi-resolution compressed sensing reconstruction via approximate message passing," *IEEE Trans. Comput. Imag.*, vol. 2, no. 3, pp. 218–234, Sep. 2016.
- [5] A. Gonzalez, H. Jiang, G. Huang, and L. Jacques, "Multi-resolution compressive sensing reconstruction," Feb. 2016, *arXiv:1602.05941*.
- [6] ERGO, "European robotic goal-oriented autonomous controller (ERGO)." Accessed: Jun. 2, 2020. [Online]. Available: <https://www.h2020-ergo.eu>
- [7] P. Bhattacharjee and A. Börner, "A plug-n-play framework and acquisition methodology for remote exploration systems with single pixel cameras," in *Proc. SPIE 11062, Digit. Opt. Technol.*, Munich, Germany, 2019, pp. 228–243.
- [8] E. J. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [9] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [10] E. J. Candes and J. K. Romberg, "Signal recovery from random projections," in *Proc. SPIE 5674, Comput. Imag. III*, San Jose, CA, USA, 2005, pp. 76–86.
- [11] Y. Tsaig and D. L. Donoho, "Extensions of compressed sensing," *Signal Process.*, vol. 86, no. 3, pp. 549–571, Mar. 2006.
- [12] B. Roman, A. Hansen, and B. Adcock, "On asymptotic structure in compressed sensing," Jul. 2014, *arXiv:1406.4178v2*.
- [13] B. Adcock, A. C. Hansen, C. Poon, and B. Roman, "Breaking the coherence barrier: A new theory for compressed sensing," *Forum Math., Sigma*, vol. 5, 2017, Art. no. e4.
- [14] NASA, "Mars Exploration Rovers Overview." Accessed: Jun. 1, 2020. [Online]. Available: <https://mars.nasa.gov/mer/mission/overview>
- [15] S. Maurice *et al.*, "The ChemCam instrument suite on the mars science laboratory (MSL) rover: Science objectives and mast unit description," *Space Sci. Rev.*, vol. 170, pp. 95–166, Sep. 2012.
- [16] R. Francis *et al.*, "AEGIS autonomous targeting for ChemCam on MARS science laboratory: Deployment and results of initial science team use," *Sci. Robot.*, vol. 2, no. 7, Jun. 2017, Art. no. ean4582.
- [17] M. Burger, "Bregman distances in inverse problems and partial differential equations," in *Advances in Mathematical Modeling, Optimization and Optimal Control*, Switzerland: Springer, 2016, pp. 3–33.
- [18] M. Benning and M. Burger, "Modern regularization methods for inverse problems," *Acta Numerica*, vol. 27, pp. 1–111, May 2018.
- [19] M. Burger, E. Resmerita, and L. He, "Error estimation for Bregman iterations and inverse scale space methods in image restoration," *Computing*, vol. 81, pp. 109–135, Nov. 2007.
- [20] E. Lorenz, S. Mitchell, T. Säuberlich, C. Paproth, W. Halle, and O. Frauenberger, "Remote sensing of high temperature events by the Fire-Bird mission," *Int. Arch. Photogrammetry Remote Sens. Spatial Inf. Sci.*, vol. XL-7/W3, pp. 461–467, Apr. 2015.
- [21] M. C. Kemp, P. F. Taday, B. E. Cole, J. A. Cluff, A. J. Fitzgerald, and W. R. Tribe, "Security applications of terahertz technology," in *Proc. SPIE 5070, Terahz. Mil. Secur. Appl.*, Orlando, FL, USA, 2003, pp. 44–52.
- [22] E. Kopp *et al.*, "HP3-RAD: A compact radiometer design with on-site calibration for in-situ exploration," in *Proc. SPIE 9973, Infrared Remote Sens. Instrum. XXIV*, San Diego, CA, USA, 2016, pp. 249–254.
- [23] M. Grott *et al.*, "The MASCOT radiometer MARA for the Hayabusa 2 mission," *Space Sci. Rev.*, vol. 208, no. 1, pp. 413–431, Aug. 2016.

- [24] S. Augustin, J. Hieronymus, P. Jung, and H.-W. Hübers, "Compressed sensing in a fully non-mechanical 350 GHz imaging setting," *J. Infrared Millimeter Terahertz Waves*, vol. 36, no. 5, pp. 496–512, May 2015.
- [25] M. F. Duarte *et al.*, "Single-pixel imaging via compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 83–91, Mar. 2008.
- [26] B. Lee, "Introduction to +/- 12 degree orthogonal digital micromirror devices (DMDs)," *Texas Instrum.*, Dallas, TX, USA, 2008, Accessed: Jun. 2, 2020. [Online]. Available: <https://www.ti.com/lit/an/dlpa008b/dlpa008b.pdf>
- [27] N. Ahmed and K. R. Rao, *Orthogonal Transforms for Digital Signal Processing*. Berlin, Germany: Springer, 1975.
- [28] E. Candès and J. Romberg, "Sparsity and incoherence in compressive sampling," *Inverse Problems*, vol. 23, no. 3, pp. 969–985, Apr. 2007.
- [29] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magn. Reson. Med.*, vol. 58, no. 6, pp. 1182–1195, Oct. 2007.
- [30] A. Majumdar and R. K. Ward, "On the choice of compressed sensing priors and sparsifying transforms for MR image reconstruction: An experimental study," *Signal Process. Image Commun.*, vol. 27, no. 9, pp. 1035–1048, Oct. 2012.
- [31] E. Candès, Y. Eldar, D. Needell, and P. Randall, "Compressed sensing with coherent and redundant dictionaries," *Appl. Comput. Harmon. Anal.*, vol. 31, pp. 59–73, Jul. 2011.
- [32] S. R. Becker, E. J. Candes, and M. C. Grant, "Templates for convex cone problems with applications to sparse signal recovery," *Math. Program. Comput.*, vol. 3, no. 3, pp. 165–218, Aug. 2011.
- [33] S. Boyd, J. Duchi, and L. Vandenberghe, "Subgradients." 2018. [Online]. Available: https://web.stanford.edu/class/ee364b/lectures/subgradients_notes.pdf
- [34] D. P. Bertsekas, *Convex Optimization Theory*. Belmont, MA, USA: Athena Sci., 2009.
- [35] NASA, "Airborne visible infrared spectrometer (AVIRIS)," NASA Jet Propulsion Laboratory, Accessed: Jun. 2, 2020. [Online]. Available: <https://aviris.jpl.nasa.gov/>
- [36] SPECTIR, "Free data samples." *SPECTIR Adv. hyperspectral solutions*. Accessed: Jun. 2, 2020, [Online]. Available: <http://www.spectir.com/free-data-samples/>
- [37] A. A. Mueller, A. Hausold, and P. Strobl, "HySens-DAIS/ROSIS imaging spectrometers at DLR," in *Proc. SPIE 4545, Remote Sens. Environ. Monit. GIS Appl. Geol.*, Toulouse, France, 2002, pp. 225–235.
- [38] MSL Mars CHEMCAM Libs Spectra 4/5 RDR V1.0, PDS Geoscience Node, Oct. 2019, doi: [10.17189/1519485](https://doi.org/10.17189/1519485).
- [39] I. Daubechies, *Ten Lectures on Wavelets*. Philadelphia, PA, USA: SIAM, 1992.
- [40] E. van den Berg and M. P. Friedlander, "SPOT-A linear operator toolbox for Matlab." 2013. [Online]. Available: <https://github.com/mpf/spot>
- [41] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [42] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*. New York, NY, USA: Birkhauser, 2013.
- [43] M. Benning and M. Burger, "Error estimates for general fidelities," *Electron. Trans. Numer. Anal.*, vol. 38, pp. 44–68, Mar. 2011.
- [44] X. Wang, "Multi-resolution approximate message passing." 2016. [Online]. Available: <https://github.com/xingwangsfu/MR-AMP/>
- [45] T. Goldstein, "Implementation of the STOne transform for compressed sensing video." 2018. [Online]. Available: <https://github.com/tomgoldstein/stone/>
- [46] N. Vaswani and W. Lu, "Modified-CS: Modifying compressive sensing for problems with partially known support," *IEEE Trans. Signal Process.*, vol. 58, no. 9, pp. 4595–4607, Sep. 2010.



Protim Bhattacharjee was born in New Delhi, India, in 1991. He studied electronics and communications engineering from Indraprastha Institute of Information Technology Delhi, New Delhi, India. He is currently working toward the Doctoral degree with the Department of Electrical Engineering, Friedrich-Alexander-University Erlangen-Nürnberg, Erlangen, Germany. He is also with the German Aerospace Center, Berlin, Germany. His research interests include computational imaging, statistical estimation, and signal processing methods for remote sensing.



Martin Burger was born in Wels, Austria, in 1976. He studied mathematics from Johannes Kepler University Linz, Linz, Austria, and the University of Milano, Milan, Italy, and received the Ph.D. degree in 2000. After working as a Postdoc with Johannes Kepler University, CAM Assistant Professor with the University of California, Los Angeles, CA, USA, and a Scientific Adviser of the Johann Radon Institute for Computational and Applied Mathematics, he accepted an offer from the Westfälische-Wilhelms University, Münster, Germany, where he was a Full Professor from 2006 to 2018. Since 2018, he has been a Full Professor in applied mathematics with Friedrich-Alexander-University Erlangen-Nürnberg, Erlangen, Germany. His research interests include inverse problems, mathematical imaging, and theoretical foundations of deep learning. Dr. Burger was the recipient of the numerous awards and honours for his work in applied mathematics and inverse problems, such as the Calderon Prize of the International Problems International Association and an ERC Consolidator Grant in 2013. He is on the Editorial Board of several journals in applied mathematics, including *Inverse Problems*, *Inverse Problems and Imaging*, and *Mathematics of Computations*. Since 2017, he has been the Editor-in-Chief of the *European Journal of Applied Mathematics*.



Anko Börner was born in Berlin, Germany. He studied electrical engineering from Technical University Ilmenau, Ilmenau, Germany, and received the Doctoral degree in 1999. In 1996, he joined the German Aerospace Center (DLR) for the Ph.D. studies on data processing on-board satellites. In 1996, he was a Postdoc with Zurich University, Switzerland. In 2000, he came back to Berlin to become a Scientific Researcher with the German Aerospace Center. Since 2003, he has been the Head of the department in different institutes. His research interests include modeling and simulation of optical systems, computer vision and sensor artificial intelligence. He was involved in several ESA and NASA space missions. In 2011, Dr. Börner was awarded with the DLR research grant, which allows him to do a research stay at Auckland, New Zealand. In 2018, he and his team won the Innovation Award Berlin/Brandenburg for developing an optical GPS-free navigation system.



Veniamin I. Morgenshtern was born in Leningrad, Russia, in 1982. He received the Specialist degree in mathematics and software engineering from Saint-Petersburg State University, Saint Petersburg, Russia, in 2004, and the Ph.D. degree in electrical engineering from ETH Zurich, Zurich, Switzerland, in 2010. From 2010 to 2012, he was a Postdoctoral Researcher with Electrical Engineering Department, ETH Zurich. From 2012 to 2016, he was a Postdoctoral Researcher with Statistics Department, Stanford University, Stanford, CA, USA. In 2017, he was a Researcher with Helm. AI, a self-driving car startup. From 2017 to 2018, he was the Chief Scientist with Mentality.AI, an algorithmic consulting company. Since 2018, he has been a Professor in machine learning and signal processing with Friedrich-Alexander-University Erlangen-Nürnberg (FAU), Erlangen, Germany. His research interests include mathematical signal processing, statistical machine learning, and deep learning. Dr. Morgenshtern was the recipient of the ETH Zurich Medal for outstanding doctoral thesis in 2011, and of the Fellowship for Advanced Researchers from the Swiss National Science Foundation in 2012. He was the recipient of Second Prize in Thomson Reuters Eikon Text Tagging Challenge, a machine learning and natural language processing competition, and the teaching award for the Best Laboratory Course at the Technical Faculty at FAU.