

Fast Integral Image Estimation at 1% measurement rate

Kuldeep Kulkarni, Pavan Turaga

Abstract—We propose a framework called **ReFIInE** to directly obtain integral image estimates from a very small number of spatially multiplexed measurements of the scene without iterative reconstruction of any auxiliary image, and demonstrate their practical utility in visual object tracking. Specifically, we design measurement matrices which are tailored to facilitate extremely fast estimation of the integral image, by using a single-shot linear operation on the measured vector. Leveraging a prior model for the images, we formulate a nuclear norm minimization problem with second order conic constraints to jointly obtain the measurement matrix and the linear operator. Through qualitative and quantitative experiments, we show that high quality integral image estimates can be obtained using our framework at very low measurement rates. Further, on a standard dataset of 50 videos, we present object tracking results which are comparable to the state-of-the-art methods, even at an extremely low measurement rate of 1%.

Index Terms—Spatial-multiplexing cameras, Nuclear-norm minimization, Low sensing rates, Tracking, Integral Images



1 INTRODUCTION

In this paper, we study the problem of obtaining integral image estimates from emerging flexible programmable imaging devices. These novel imaging devices, often considered under the broad umbrella of spatial-multiplexing cameras (SMCs)[1], [2], [3] provide a number of benefits in reducing the amount of sensing for portable and resource constrained acquisition. The imaging architectures in these cameras employ spatial light modulators like digital micromirror arrays to optically compute projections of the scene. Mathematically, the projections are given by $y = \phi x$, where $x \in \mathbb{R}^n$ is the image, $y \in \mathbb{R}^m$, known as the measurement vector, denotes the set of sensed projections and $\phi \in \mathbb{R}^{m \times n}$ is called measurement matrix defined by the set of multiplexing patterns. The nature of the acquisition framework enables us to deploy SMCs in resource constrained settings, wherein one can employ $m \ll n$ number of photon detectors to sense otherwise high-resolution imagery [2], [4] and obtain a very small number of measurements. Later, a reconstruction algorithm is used to recover the image x . However, reconstructing x from y when $m < n$ is an ill-posed problem. Researchers in the past have attempted to provide solutions by carefully designing the measurement matrix ϕ in the hope of easier recovery of x from y . Recent compressive sensing (CS) theory provides one possible solution to tackle the above mentioned ill-posed problem. According to CS theory, a signal can be recovered perfectly from a small number of $m = \mathcal{O}(s \log(\frac{n}{s}))$ such pseudo-random (PR)

multiplexed measurements, where s is the sparsity of the signal. However, a significant research shows that high-quality reconstruction is computationally intensive [5], [6], [7], [8]. Hence, despite the promise of CS-based SMCs [2], [3], the computational bottleneck of non-linear, iterative reconstruction has withheld their wide-spread adoption in applications which require fast inference of objects, actions, and scenes. This has led to researchers exploring the option of tackling inference problems directly from these pseudo-random multiplexed measurements [4], [9], [10], [11], [12], [13], [14] (more on these later in the section in related work). However, the ‘universal’ nature of such measurements has made it challenging to devise new or adopt existing computer vision algorithms to solve the inference problem at hand.

The need to acquire as less data as possible combined with the limitations of pseudo-random multiplexers props us to outline the following goal. The goal of this paper is to propose a novel sensing framework for SMCs such that acquired measurements satisfy the following properties. 1) The measurements are not random in nature but are tailored for a particular application. 2) The number of measurements is 2 orders less than the number of pixels, so that SMCs based on our framework can be employed in resource constrained applications. 3) A simple linear operation on the measurement vector y yields a ‘proxy’ representation (e.g integral images, gradient images) from which the required features are extracted for the application in hand, thus avoiding the computationally expensive iterative and non-linear reconstruction.

In this paper, we focus on one such ‘proxy’ representation, integral images. Integral images are extremely attractive representation since Haar-like features and box-filtered image outputs can be computed from integral images with a small and fixed number of floating point operations in constant time [15]. These advantages have led to their

• K. Kulkarni and P. Turaga are with the School of Arts, Media and Engineering and School of Electrical, Computer and Energy Engineering, Arizona State University. Email: kkulkar1@asu.edu, pturaga@asu.edu.

widespread use in real time applications like face detection [15], pedestrian detection [16], object tracking [17], [18], [19], [20] and object segmentation [21].

Instead of setting a fixed number of measurements, we formulate an optimization problem to minimize the number of measurements while incorporating the other two (1 and 3) properties of measurements (as mentioned above) in the constraints. Minimizing the number of measurements is akin to minimizing the rank of the measurement matrix. In more concrete terms, the problem is posed to jointly minimize the rank of the measurement matrix, ϕ and learn the linear operator, \mathcal{L} which when applied on the measurement vector yields the approximate integral image, with the probabilistic constraint that the error between the approximate integral image and the exact integral image is within allowable limits with high probability. By controlling the allowable error limit, we can obtain measurement matrix of the desired rank. Incorporating a wavelet domain prior model for natural images combined with a relaxation (explained in section 2) allows us to convert the probabilistic constraint into a series of second conic constraints. Rank minimization is a NP-hard problem. Relaxing the objective function to nuclear norm allows to use off-the-shelf convex optimization tools to solve the problem and obtain the measurement matrix and the linear operator.

Contributions: 1) The main contribution of the paper is a novel framework to recover estimates of integral images from a small number of spatially multiplexed measurements without iterative reconstruction of any auxiliary image. We dub the framework **ReFINE** (Reconstruction-Free Integral Image Estimation). 2) Leveraging the MGGD (multivariate generalized Gaussian distribution) prior model for the vector of detailed wavelet coefficients of natural images, we propose a nuclear norm minimization formulation to obtain a new specialized measurement matrix. We term the measurements acquired with such a measurement matrix, as **ReFINE** measurements. 3) On a large dataset of 4952 images, we present qualitative and quantitative results to show that high quality estimates of integral images and box-filtered outputs can be recovered from **ReFINE** measurements in real-time. 4) We show object tracking results, which are comparable to state-of-the-art methods, on a challenging dataset of 50 videos to demonstrate the utility of the box-filtered output estimates in tackling inference problems from SMCs at 1% measurement rate.

Related Work: The related previous works in literature follow one of the two themes. Some attempt to tackle inference problems directly from PR measurements without optimizing for the measurement matrix for the inference task at hand, some others attempt to optimize measurement matrix for a particular signal model so as to minimize reconstruction error.

a) Design of measurement matrix: A closely related work can be found in [22], wherein a framework is proposed to jointly optimize for a measurement matrix and an overcomplete sparsifying dictionary for small patches of images. Results suggest that better reconstruction results can be obtained using this strategy. However, learning

global dictionaries for entire images is not possible, and hence the framework is not scalable. Goldstein *et al.*[23] designed measurement matrices called ‘STOne’ Transform which facilitate fast low resolution ‘previews’ just by direct reconstruction, and the hope is that ‘previews’ are of high enough quality so that conventional methods for inference tasks can be applied. Assuming a multi-resolutional signal model for natural images, Chang *et al.*[24] proposed an algorithm to obtain measurements which have the maximum mutual information with natural images.

b) Inference problems from CS videos: A LDS (Linear Dynamical System) based approach was proposed by Sankaranarayanan *et al.*[4] to model CS videos and recover the LDS parameters directly from PR measurements. However, the method is sensitive to spatial and view transforms. Calderbank *et al.*[12] theoretically proved that one can learn classifiers directly from PR measurements, and that with high probability the performance of the linear kernel support vector machine (SVM) classifier operating on the CS measurements is similar to that of the the best linear threshold classifier operating on the original data. A reconstruction-free framework was proposed by Thirumalai *et al.*[9] to compute optical flow based on correlation estimation between two images, directly from PR measurements. Davenport *et al.*[25] proposed a measurement domain based correlation filter approach for target classification. Here, the trained filters are first projected onto PR patterns to obtain ‘smashed filters’, and then the PR measurements of the test examples are correlated with these smashed filters. Recently, Kulkarni *et al.*[13] and Lohit *et al.*[14] extended the ‘smashed filter’ approach to action recognition and face recognition respectively, and demonstrated the feasibility and scalability of tackling difficult inference tasks in computer vision directly from PR measurements.

2 BACKGROUND

In this section, we provide a brief background on the probability model for natural images, which we rely on in the paper, and introduce notations required to set up the optimization problem to derive measurement matrix and above referred linear operator.

Probability Model of natural images: There is a rich body of literature which deals with statistical modeling of natural images. We refer to some works which are related to the probability model we use in the paper. Many successful probability models for wavelet coefficients fall under the broad umbrella of Gaussian scale mixtures (GSM) [26], [27]. Typically the coefficient space is partitioned into overlapping blocks, and each block is modeled independently as a GSM, which captures the local dependencies. This implicitly gives rise to a global model of the wavelet coefficients. Building on this framework, Lyu *et al.*[28] proposed a field of Gaussian scale mixtures (FoGSM) to explicitly model the subbands of wavelet coefficients, while treating each subband independently. However, incorporating such a general model for wavelet coefficient vector makes it

very difficult to compute the distribution for even simple functions like a linear function of the wavelet coefficient vector. Therefore, it is vital to assume a prior model which can lead to tractable computation of the distribution. It is well-known that marginal distributions of detailed wavelet coefficients follow generalized Gaussian distribution [29]. We extend this notion to multi-dimensions and model the vector of detailed wavelet coefficients by multivariate generalized Gaussian distribution (MGGD).

To put it formally, let $\mathbf{U}^T \in \mathbb{R}^{n \times n}$ be the orthogonal matrix representing the $\log_2(n)$ level wavelet transform, so that $x = \mathbf{U}w$, where $w \in \mathbb{R}^n$ is the corresponding wavelet coefficient vector. Without loss of generality, we assume that all entries in the first row of \mathbf{U}^T are $1/\sqrt{n}$ so that the first entry in w corresponds to \sqrt{n} times the mean of all entries in x . Also we denote the rest $n - 1$ rows in \mathbf{U}^T by $\mathbf{U}_{2:n}^T$. Now we can write $w = [\sqrt{n}\bar{x}, w_d]$, where \bar{x} is the mean of x and w_d is the vector of detailed coefficients. As explained above, the probability distribution of w_d (MGGD) is given by

$$f(w) = K |\Sigma_{w_d}|^{-0.5} \exp(-(w_d^T \Sigma_{w_d}^{-1} w_d)^\beta), \quad (1)$$

where Σ_{w_d} , commonly known as the scatter matrix, is equal to $\text{rank}(\Sigma_{w_d}) \Gamma((2 + n - 2)/2\beta) / \Gamma((2 + n)/2\beta)$ times the covariance matrix of w_d , $\beta \in (0, 1]$, and K is a normalizing constant. For $\beta = 1$, we obtain the probability distribution for the well-known multivariate Gaussian distribution. In the following we briefly provide a background regarding the multivariate generalized Gaussian distribution.

A linear transformation of the multivariate generalized Gaussian random vector is also a multivariate generalized Gaussian random vector.

Proposition 1: [30] Let u be the a $n \times 1$ multivariate generalized Gaussian random vector with mean $\mu_u \in \mathbb{R}^n$, and scatter matrix, $\Sigma_u \in \mathbb{R}^{n \times n}$. Let A be a $l \times n$ full rank matrix. Then the $l \times 1$ random vector, $v = Au$ has the multivariate generalized Gaussian distribution with mean $\mu_v = A\mu_u \in \mathbb{R}^l$, and scatter matrix, $\Sigma_v = A\Sigma_u A^T \in \mathbb{R}^{l \times l}$.

If v is a univariate generalized Gaussian random variable, then the probability of v falling in the range of $[\delta - \mu_v, \delta + \mu_v]$, for $\delta \geq 0$, can be found in terms of lower incomplete gamma function.

Proposition 2: If v is a univariate generalized Gaussian random variable with mean $\mu_v \in \mathbb{R}$ and scatter matrix, $\Sigma_v \in \mathbb{R}$, then the probability of v falling in the range of $[-\delta + \mu_v, \delta + \mu_v]$, for $\delta \geq 0$, is given by,

$$\mathbb{P}(|v - \mu_v| \leq \delta) = 2\gamma \left(\frac{1}{2\beta}, \left(\frac{\delta}{\Sigma_v} \sqrt{\frac{\Gamma(\frac{3}{2\beta})}{\Gamma(\frac{1}{2\beta})}} \right)^{2\beta} \right), \quad (2)$$

where $\gamma(\cdot, \cdot)$ is the lower incomplete gamma function, and $\Gamma(\cdot)$ is the ordinary gamma function.

Preliminaries: Let $\mathbf{H} \in \mathbb{R}^{n \times n}$ be the block Toeplitz matrix representing the integral operation so that the integral image, $I = \mathbf{H}x \in \mathbb{R}^n$, and h_i^T for $i = 1, \dots, n$ be the rows of \mathbf{H} . Hence, I_i equals $h_i^T x$. We wish to recover the

approximate integral image, \hat{I} from the measured vector $y = \phi x$, just by applying a linear operator \mathcal{L} on y , so that $\hat{I} = \mathcal{L}y$. For reasons which will be apparent soon, we assume $\mathcal{L} = \mathbf{H}(\phi^d)^T$, where $\phi^d \in \mathbb{R}^{m \times n}$ such that $\text{rank}(\phi^d) = \text{rank}(\phi)$. We call ϕ^d as the dual of ϕ . Thus by construction $\mathcal{L} \in \mathbb{R}^{n \times m}$. The value at location i in the approximate integral image is given by $\hat{I}_i = h_i^T (\phi^d)^T \phi x$. The distortion in integral image at location i is given by $d_i = \hat{I}_i - I_i = h_i^T ((\phi^d)^T \phi x - x)$. Noting $\mathbf{Q} = (\phi^d)^T \phi$, and $n \times n$ identity matrix by \mathbf{I} , the distortions can be compactly written as $d_i = h_i^T (\mathbf{Q} - \mathbf{I})x$. We call $\mathbf{d} = [d_1, \dots, d_n]$ as distortion vector.

3 OPTIMIZATION PROBLEM

Our aim is to search for a measurement matrix ϕ of minimum rank such that distortions, d_i are within allowable limits for all i , jointly, with high probability. \mathbf{Q} by construction is the product of two matrices of identical ranks, ϕ and $(\phi^d)^T$. Hence, we have the relation, $\text{rank}(\mathbf{Q}) = \text{rank}(\phi)$. Inspired by the phase-lifting technique used in [31] and [32], instead of minimizing the rank of ϕ , we minimize the rank of \mathbf{Q} . Now, we can formally state the optimization problem as follows.

$$\underset{\mathbf{Q}}{\text{minimize}} \quad \text{rank}(\mathbf{Q}) \quad (3)$$

$$\text{s.t.} \quad \mathbb{P}(|d_1| \leq \delta_1, \dots, |d_i| \leq \delta_i, \dots, |d_n| \leq \delta_n) \geq 1 - \epsilon,$$

where $\delta_i \geq 0$ denotes the allowable limit of distortion at location i of integral image, and $0 < \epsilon < 1$. Once \mathbf{Q}^* is found, we show later in the section that the SVD decomposition of \mathbf{Q}^* allows us to write \mathbf{Q}^* as a product of two matrices of identical ranks, thus yielding both the measurement matrix, ϕ^* and the desired linear operator, \mathcal{L}^* . The constraint in (3) is a probabilistic one. Hence to compute it, one needs to assume a statistical prior model for x . Using the model in 1, and its properties given in proposition (1) and (2), we arrive at a solvable optimization problem.

Computation of probabilistic constraint in (3): Substituting for x , we can write the distortion at location i as $d_i = h_i^T (\mathbf{Q} - \mathbf{I})\mathbf{U}w$. We let all the entries in the first row of ϕ and ϕ^d to be equal to $1/\sqrt{n}$, so that one of the m measurements is exactly equal to $\sqrt{n}\bar{x}$. Further, we denote the rest $m - 1$ rows of the two matrices by $\phi_{2:m}$ and $\phi_{2:m}^d$. Now, if we restrict $\phi_{2:m}$ and $\phi_{2:m}^d$ to be respectively equal to $\mathbf{C}\mathbf{U}_{2:n}^T$ and $\mathbf{D}\mathbf{U}_{2:n}^T$ for some \mathbf{C}, \mathbf{D} in $\mathbb{R}^{(m-1) \times (n-1)}$, then from basic linear algebra we can show that $d_i = h_i^T (\mathbf{P} - \mathbf{I})\mathbf{U}_{2:n}^T w_d$, where $\mathbf{P} = (\phi_{2:m}^d)^T \phi_{2:m}$. It is easy to see that $\mathbf{Q} = \mathbf{P} + \frac{1}{n}\mathbf{O}$, and $\text{rank}(\mathbf{Q}) = \text{rank}(\mathbf{P}) + 1$, where \mathbf{O} is the matrix with all its entries equal to unity (see Appendix A for details). Hence we can replace the objective function in (3) by $\text{rank}(\mathbf{P})$. Rank minimization is a non-convex problem. Hence we relax the objective function to nuclear norm, as is done typically. To compute the constraint in (3), one needs to first compute the joint probability of $\mathbf{d} = [d_1, \dots, d_n]$, and then compute a n dimensional definite integral. Now that d_i 's are linear

combinations of w_d , it follows from proposition 1, that \mathbf{d} also has a multivariate generalized Gaussian distribution. However, no closed form for the definite integral is known. Hence, we relax the constraint by decoupling it into n independent constraints, each enforcing the constraint that the distortion at a specific location is to be within allowable limits with high probability, independent of the distortions at other locations. The optimization with relaxed constraints is thus given by

$$\begin{aligned} & \underset{\mathbf{P}}{\text{minimize}} \quad \|\mathbf{P}\|_* \\ & \text{s.t} \quad \mathbb{P}(|d_i| \leq \delta_i) \geq 1 - \epsilon, \quad i = 1, \dots, n. \end{aligned} \quad (4)$$

Now, $d_i = h_i^T (\mathbf{P} - \mathbf{I}) \mathbf{U}_{2:n}^T w_d$, is a linear combination of the entries of w_d . From the proposition 2, d_i has a one-dimensional generalized Gaussian distribution with zero mean and scatter parameter, $\left\| \Sigma_{w_d}^{1/2} \mathbf{U}_{2:n}^T (\mathbf{P} - \mathbf{I})^T h_i \right\|$, and the probability in equation 4 can be explicitly written as follows.

$$\begin{aligned} & \mathbb{P}(|d_i| \leq \delta_i) \\ & = 2\gamma \left(\frac{1}{2\beta}, \left(\frac{\delta_i}{\left\| \Sigma_{w_d}^{1/2} \mathbf{U}_{2:n}^T \mathbf{P}^T h_i - \Sigma_{w_d}^{1/2} \mathbf{U}_{2:n}^T h_i \right\|_2} \sqrt{\frac{\Gamma(\frac{3}{2\beta})}{\Gamma(\frac{1}{2\beta})}} \right)^{2\beta} \right). \end{aligned} \quad (5)$$

The optimization problem now can be rewritten as

$$\begin{aligned} & \underset{\mathbf{P}}{\text{minimize}} \quad \|\mathbf{P}\|_* \quad \text{s.t} \\ & 2\gamma \left(\frac{1}{2\beta}, \left(\frac{\delta_i}{\left\| \Sigma_{w_d}^{1/2} \mathbf{U}_{2:n}^T \mathbf{P}^T h_i - \Sigma_{w_d}^{1/2} \mathbf{U}_{2:n}^T h_i \right\|_2} \sqrt{\frac{\Gamma(\frac{3}{2\beta})}{\Gamma(\frac{1}{2\beta})}} \right)^{2\beta} \right) \geq 1 - \epsilon \\ & \left. \begin{aligned} & \mathbf{P}^k = \arg \min_{\mathbf{P}} \tau \|\mathbf{P}\|_* + \frac{1}{2} \|\mathbf{P}\|_F^2 + \langle \mathbf{z}^{k-1}, \mathcal{F}(\mathbf{P}) \rangle \\ & \mathbf{z}_i^k = P_i(\mathbf{z}_i^{k-1} + \eta^k f_i(\mathbf{P}^k)), \quad i = 1, \dots, n \end{aligned} \right\} \end{aligned} \quad (10)$$

We compactly write $\Sigma_{w_d}^{1/2} \mathbf{U}_{2:n}^T \mathbf{P}^T h_i$ as $\mathcal{A}_i(\mathbf{P})$, $\Sigma_{w_d}^{1/2} \mathbf{U}_{2:n}^T h_i$ as b_i and $\frac{\delta_i}{(\gamma^{-1}(\frac{1}{2\beta}, \frac{1-\epsilon}{2}))^{\frac{1}{2\beta}}} \sqrt{\frac{\Gamma(\frac{3}{2\beta})}{\Gamma(\frac{1}{2\beta})}}$ as Δ_i . Plugging above in equation (5), and rearranging terms, we have

$$\begin{aligned} & \underset{\mathbf{P}}{\text{minimize}} \quad \|\mathbf{P}\|_* \\ & \text{s.t} \quad \|\mathcal{A}_i(\mathbf{P}) - b_i\|_2 \leq \Delta_i \quad i = 1, \dots, n. \end{aligned} \quad (6)$$

We can rewrite the problem in conic form as below.

$$\begin{aligned} & (P1) \quad \underset{\mathbf{P}}{\text{minimize}} \quad \|\mathbf{P}\|_* \quad \text{s.t} \\ & \left[\begin{array}{c} b_i - \mathcal{A}_i(\mathbf{P}) \\ \Delta_i \end{array} \right] \in \mathcal{K}_i, \quad i = 1, \dots, n, \end{aligned} \quad (7)$$

where \mathcal{K}_i is a second order cone $\mathcal{K}_i = \{(x_i, t_i) \in \mathbb{R}^{n+1} : \|x_i\| \leq t_i\}$. Let $\mathbf{b} = [b_1, \dots, b_n]$, and $\mathcal{A}(\mathbf{P}) = [\mathcal{A}_1(\mathbf{P}), \dots, \mathcal{A}_n(\mathbf{P})]$. Let \mathcal{A}^* denote the adjoint of the linear operator \mathcal{A} . It is easy to recognize that the optimization above is a convex problem, since nuclear norm is convex and the constraints enforce finite bounds on the norms of affine functions of the decision variable, \mathbf{P} and hence are also convex. Even though the constraints are second-order cone constraints, the standard second order conic programming methods cannot be used to solve (P1)

since nuclear norm is non-smooth. The nuclear norm is smoothed by the addition of a square of Frobenius norm of the matrix, and is replaced by $\tau \|\mathbf{P}\|_* + \frac{1}{2} \|\mathbf{P}\|_F^2$, where $\tau > 0$. The optimization problem with the smoothed objective function is given in 8.

$$\begin{aligned} & (P2) \quad \underset{\mathbf{P}}{\text{minimize}} \quad \tau \|\mathbf{P}\|_* + \frac{1}{2} \|\mathbf{P}\|_F^2 \\ & \text{s.t} \quad \left[\begin{array}{c} b_i - \mathcal{A}_i(\mathbf{P}) \\ \Delta_i \end{array} \right] \in \mathcal{K}_i, \quad i = 1, \dots, n. \end{aligned} \quad (8)$$

Recently, many algorithms [33], [34] have been developed to tackle nuclear norm minimization problem of this form in the context of matrix completion. We use SVT (singular value thresholding) algorithm [34] to solve (P2).

SVT iteration to solve (P2): Here, we first briefly describe the SVT algorithm for smoothed nuclear norm minimization with general convex constraints, and later we show how we adapt the same to our problem P2 which has n second-order constraints. Let the smoothed nuclear norm with general convex constraints, be given as below.

$$\begin{aligned} & \underset{\mathbf{P}}{\text{minimize}} \quad \tau \|\mathbf{P}\|_* + \frac{1}{2} \|\mathbf{P}\|_F^2 \\ & \text{s.t} \quad f_i(\mathbf{P}) \leq 0, \quad i = 1, \dots, n, \end{aligned} \quad (9)$$

where $f_i(\mathbf{P}) \leq 0$, $i = 1, \dots, n$ denote the n convex constraints. Let $\mathcal{F}(\mathbf{P}) = [f_1(\mathbf{P}), \dots, f_n(\mathbf{P})]$. The SVT algorithm for the 9 with the modified objective function is given as below.

where \mathbf{z}^k is a short form for $[\mathbf{z}_1^k, \dots, \mathbf{z}_n^k]$, and $P_i(\mathbf{q})$ denotes the projection of \mathbf{q} onto the convex set defined by the constraint $f_i(\mathbf{P}) \leq 0$. Let $\mathbf{z}_i^k = [\mathbf{y}_i^k, s_i^k]$, so that the vector \mathbf{y}_i^k denotes the first n elements of \mathbf{z}_i^k and s_i^k denotes the last element of \mathbf{z}_i^k . Let \mathbf{y}^k be a short form for $[\mathbf{y}_1^k, \dots, \mathbf{y}_n^k]$, and s^k is a short form for $[s_1^k, \dots, s_n^k]$. To obtain an explicit form of the update equations in 10 for our problem, P1, let us consider the first equation of the same. $\mathcal{F}(\mathbf{P})$ for P1 is given by $[b_1 - \mathcal{A}_1(\mathbf{P}), \Delta_1, \dots, b_n - \mathcal{A}_n(\mathbf{P}), \Delta_n]^T$. We substitute for $\mathcal{F}(\mathbf{P})$, and after removal of the terms not involving \mathbf{P} , we have

$$\mathbf{P}^k = \arg \min_{\mathbf{P}} \tau \|\mathbf{P}\|_* + \frac{1}{2} \|\mathbf{P}\|_F^2 + \langle \mathbf{y}^{k-1}, \mathbf{b} - \mathcal{A}(\mathbf{P}) \rangle. \quad (11)$$

Equation 11 can be rewritten as below.

$$\begin{aligned} \mathbf{P}^k = \arg \min_{\mathbf{P}} \quad & \tau \|\mathbf{P}\|_* + \frac{1}{2} \|\mathbf{P} - \mathcal{A}^*(\mathbf{y}^{k-1})\|_F^2 \\ & - \frac{1}{2} \|\mathcal{A}^*(\mathbf{y}^{k-1})\|_F^2 + \langle \mathbf{P}, \mathcal{A}^*(\mathbf{y}^{k-1}) \rangle \\ & + \langle \mathbf{y}^{k-1}, \mathbf{b} \rangle - \langle \mathbf{y}^{k-1}, \mathcal{A}(\mathbf{P}) \rangle. \end{aligned}$$

Removing the terms not involving \mathbf{P} and noting that $\langle \mathbf{P}, \mathcal{A}^*(\mathbf{y}^{k-1}) \rangle = \langle \mathbf{y}^{k-1}, \mathcal{A}(\mathbf{P}) \rangle$, we have the following.

$$\mathbf{P}^k = \arg \min_{\mathbf{P}} \tau \|\mathbf{P}\|_* + \frac{1}{2} \|\mathbf{P} - \mathcal{A}^*(\mathbf{y}^{k-1})\|_F^2. \quad (12)$$

Before we write down the solution to equation 12, we first define \mathcal{D}_τ , the singular value shrinkage operator. Consider the SVD of a matrix \mathbf{X} , given by $\mathbf{X} = \mathbf{W}\mathbf{\Sigma}\mathbf{V}^T$. Then for $\tau \geq 0$, the singular value shrinkage operator, \mathcal{D}_τ is given by $\mathcal{D}_\tau(\mathbf{X}) = \mathbf{W}\mathcal{D}_\tau(\mathbf{\Sigma})\mathbf{V}^T$, $\mathcal{D}_\tau(\mathbf{\Sigma}) = \text{diag}(\{(\sigma_i - \tau)_+\})$, where $t_+ = \max(0, t)$. The solution to equation 12 is given by $\mathbf{P}^k = \mathcal{D}_\tau(\mathcal{A}^*(\mathbf{y}^{k-1}))$. Now, it remains to calculate $\mathcal{A}^*(\mathbf{y}^{k-1})$. We achieve it according to the following. Consider $\langle \mathcal{A}^*(\mathbf{y}^{k-1}), \mathbf{P} \rangle$.

$$\begin{aligned} & \langle \mathbf{P}, \mathcal{A}^*(\mathbf{y}^{k-1}) \rangle \\ &= \langle \mathcal{A}(\mathbf{P}), \mathbf{y}^{k-1} \rangle = \sum_{i=1}^n \langle \mathcal{A}_i(\mathbf{P}), \mathbf{y}_i^{k-1} \rangle \\ &= \sum_{i=1}^n \langle \mathbf{P}, \mathcal{A}_i^*(\mathbf{y}_i^{k-1}) \rangle = \langle \mathbf{P}, \sum_{i=1}^n \mathcal{A}_i^*(\mathbf{y}_i^{k-1}) \rangle \end{aligned}$$

Hence, we have $\mathcal{A}^*(\mathbf{y}^{k-1}) = \sum_{i=1}^n \mathcal{A}_i^*(\mathbf{y}_i^{k-1})$. Thus, the first equation of SVT iteration for our problem is given by

$$\mathbf{P}^k = \mathcal{D}_\tau \left(\sum_{i=1}^n \mathcal{A}_i^*(\mathbf{y}_i^{k-1}) \right). \quad (13)$$

Using basic linear algebra, it can be shown that $\mathcal{A}_i^*(\mathbf{y}_i^{k-1}) = \mathbf{U}_{2:n}(\mathbf{\Sigma}_{w_d}^{1/2})^T \mathbf{y}_i^{k-1} h_i^T$ (see Appendix B for details).

We now provide the projection onto the convex cone \mathcal{K}_i . The projection operator, $P_{\mathcal{K}_i}$ as derived in [35] is given as follows.

$$P_{\mathcal{K}_i} : (x, t) \mapsto \begin{cases} (x, t), & \|x\| \leq t, \\ \frac{\|x\|+t}{2\|x\|}(x, \|x\|), & -\|x\| \leq t \leq \|x\|, \\ (0, 0), & t \leq -\|x\|. \end{cases} \quad (14)$$

To solve (P2), starting with $\begin{bmatrix} \mathbf{y}_i^0 \\ s_i^0 \end{bmatrix} = \mathbf{0}$ for all $i = 1, \dots, n$, the k^{th} SVT iteration is given by (15).

$$\left. \begin{aligned} \mathbf{P}^k &= \mathcal{D}_\tau \left(\sum_{i=1}^n \mathbf{U}_{2:n}(\mathbf{\Sigma}_{w_d}^{1/2})^T \mathbf{y}_i^{k-1} h_i^T \right) \\ \begin{bmatrix} \mathbf{y}_i^k \\ s_i^k \end{bmatrix} &= P_{\mathcal{K}_i} \left(\begin{bmatrix} \mathbf{y}_i^{k-1} \\ s_i^{k-1} \end{bmatrix} + \eta^k \begin{bmatrix} b_i - \mathcal{A}_i(\mathbf{P}^k) \\ -\Delta_i \end{bmatrix} \right), \end{aligned} \right\} \quad (15)$$

where, \mathcal{A}_i^* are the adjoints of linear operators \mathcal{A}_i . For the iterations (15) to converge, we need to choose the step sizes, $\eta^k \leq \frac{2}{\|\mathcal{A}\|_2^2}$, where $\|\mathcal{A}\|_2$ is the spectral norm of the linear transformation \mathcal{A} [34].

Once the solution \mathbf{P}^* is found, using the relation noted earlier in the section, we have $\mathbf{Q}^* = \mathbf{P}^* + \frac{1}{n}\mathbf{O}$. Having obtained \mathbf{Q}^* , the task now is to express it into a product of two matrices of identical ranks. This is done almost trivially as follows. Noting the singular value decomposition of \mathbf{Q}^* as $\mathbf{Q}^* = \mathbf{W}_M \mathbf{\Sigma}_M \mathbf{V}_M^T$, where $\mathbf{\Sigma}_M = \text{diag}\{\lambda_1, \dots, \lambda_M\}$ denote the diagonal matrix with the non-zero singular values arranged along its diagonal, and \mathbf{W}_M and \mathbf{V}_M^T are the matrices whose columns are the left and right singular vectors respectively. We can choose $\phi^* = \mathbf{\Sigma}_M^{1/2} \mathbf{V}_M^T$, so that $((\phi^d)^*)^T = \mathbf{W}_M \mathbf{\Sigma}_M^{1/2}$ and $\text{rank}(\phi^*) = \text{rank}((\phi^d)^*) = \text{rank}(\mathbf{Q}^*)$. We, henceforth refer to ϕ^* as **ReFINE** ϕ , and

the corresponding measurements, $y = \phi^* x$ as **ReFINE** measurements. The desired linear operator \mathcal{L}^* is given by $\mathbf{H}\mathbf{W}_M \mathbf{\Sigma}_M^{1/2}$. By construction, the approximate integral image \hat{I} is given by $\hat{I} = \mathcal{L}^* y = (\mathbf{H}\mathbf{W}_M \mathbf{\Sigma}_M^{1/2}) y$.

Computational Complexity: Since the length of the image x is n , the number of entries in \mathbf{P} is n^2 . Hence, the dimension of the optimization problem P2 is n^2 . This means that if we are to optimize for a measurement matrix to operate on an image of size, 256×256 , the dimension of the optimization problem would be 2^{32} ! Optimizing over a large number of variables is computationally expensive, if not impractical. Hence we propose the following suboptimal solution to obtain the measurement matrix. We divide the image into non-overlapping blocks of fixed size, and sense each block using a measurement matrix, optimized for this fixed size, independently of other blocks. Let the image, x be divided into B blocks, x_1, x_2, \dots, x_B , each of a fixed size, $f \times f$, and $\phi_f \in \mathcal{R}^{m \times f^2}$ be the measurement matrix optimized for images of size $f \times f$, and $(\phi_f^d)^T$ be the corresponding dual matrix. Then the ‘**ReFINE**’ measurements, y are given by the following,

$$y = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_B \end{bmatrix} = \begin{bmatrix} \phi_f & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \phi_f & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \phi_f \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_B \end{bmatrix}. \quad (16)$$

Once the measurements, y are obtained, the integral image, \hat{I} is given by

$$\hat{I} = \mathbf{H} \begin{bmatrix} (\phi_f^d)^T & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & (\phi_f^d)^T & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & (\phi_f^d)^T \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_B \end{bmatrix}. \quad (17)$$

4 EXPERIMENTS

Before we can conduct experiments to evaluate our framework, we first need to estimate the parameters of the probability model in 1. Estimating parameters of the probability model and optimizing measurement matrices for any arbitrary large sized image blocks is not practical since the former task requires an enormous amount of image data and the latter requires prohibitive amount of memory and computational resources. Hence, we fix the block size to be 32×32 images. The scatter matrix $\mathbf{\Sigma}_{w_d}$ is a scalar multiple of the covariance matrix of w_d . Hence it suffices to compute the covariance matrix. To this end, we first downsample all the 5011 training images in PASCAL VOC 2007 dataset [36] to a size of 32×32 , so that $n = 1024$ and then obtain the level 7 Daubechies wavelet coefficient vectors. We compute the sample covariance matrix of thus obtained wavelet coefficient data. For various values of β , we evaluate the χ^2 distance between the histograms of the individual wavelet coefficients and their respective theoretical marginal distributions with the variances computed above. We found for $\beta = 0.68$, the distance computed above is minimum.

Computing measurement matrix: To obtain a measurement matrix, we need to input a desired distortion vector δ to the optimization problem in (P2). The desired distortion vector is computed according to the following. We first perform principal component analysis (PCA) on the downsampled 5011 training images in the PASCAL VOC 2007 dataset [36]. We use only the top 10 PCA components as ϕ to ‘sense’ these images. We obtain the desired distortion vector by first assuming $\phi^d = \phi$ and calculating distortions, $|d_i^j|$ at each location for all training images, $j = 1, \dots, 5011$. Now, the entry in location i of the desired δ is given by the minimum value α , so that 95% of the values, $|d_i^j|, j = 1, \dots, 5011$ are less than α . We use $\epsilon = 0.95$ and solve (P2) to obtain \mathbf{P}^* , and hence also \mathbf{Q}^* . The rank of **ReFINE** ϕ is simply the rank of \mathbf{Q}^* .

Estimation of integral images: We show that good quality estimates of integral images can be obtained using our framework. To this end, we first construct **ReFINE** measurement matrices of various ranks, M . We achieve this by considering the SVD of \mathbf{Q}^* obtained above. For a particular value of M , the **ReFINE** ϕ is calculated according to $\phi_M = \Sigma_M^{1/2} \mathbf{V}_M^T$, where $\Sigma_M = \text{diag}\{\lambda_1, \dots, \lambda_M\}$ is a diagonal matrix with M largest singular values arranged along the diagonal and \mathbf{V}_M^T denote the corresponding rows in \mathbf{V}^T . Its dual, ϕ^d , is calculated similarly. For each particular measurement rate, determined by the value of M , the integral image estimates are recovered from M **ReFINE** measurements for all the 4952 test images in the PASCAL VOC 2007 dataset [36]. Similarly integral image estimates are recovered from random Gaussian measurements by first performing non-linear iterative reconstruction using the CoSamP algorithm [8] and then applying the integral operation on the reconstructed images. This pipeline is used as baseline to compare integral estimates, and henceforth is referred to as ‘RG-CoSamP’. We then measure the recovered signal-to-noise ratio (RSNR) via $20 \log_{10} \left(\frac{\|\hat{I}\|_F}{\|\hat{I} - I\|_F} \right)$. The average RSNR for recovered integral image estimates as well as the time taken to obtain integral images are tabulated in the table 1. Our framework outperforms RG-CoSamP in terms of both recovery signal-to-noise ratio and time taken to estimate the integral image, at all measurement rates. This shows that **ReFINE** ϕ , the measurement matrices designed by our framework, facilitate faster and more accurate recovery of integral image estimates than the universal matrices. The average time taken to obtain integral image estimates in our framework is about 0.003s, which amounts to a real-time speed of 300 FPS. Further, we randomly select four images (‘Two Men’, ‘Plane’, ‘Train’ and ‘Room’) from the test set (shown in figure 1(a), 1(b), 1(c), 1(d)) and present qualitative and quantitative results for individual images. Image-wise RSNR v/s measurement rate plots are shown in figure 2. It is very clear that for all the four images, our framework clearly outperforms RG-CoSamP in terms of RSNR, at all measurement rates.

Estimation of box-filtered outputs: It is well known that box-filtered outputs of any size can efficiently computed using integral images [15]. To show the capability

of our framework in recovering good quality box-filtered output estimates, we conduct the following experiment. For box filters of sizes 3×3 , 5×5 and 7×7 , we compute the estimates of filtered outputs for the four images using their respective recovered integral image estimates. RSNR v/s measurement rate plots for different filter sizes are shown in figure 3. It is evident that even for a remarkably low measurement rate of 1% , we obtain high RSNR box-filtered outputs. For a fixed measurement rate, expectedly the RSNR increases with the size of the filter. This shows the structures which are more global in nature are captured better. This is particularly true in the case of ‘Plane’ image. The high RSNR for this image hints at the absence of fine structures and homogeneous background. Further, for the ‘Two Men’ Image, we also compare the heat maps of the exact box-filtered outputs with the estimated ones. We fix the measurement rate to 1%. For filter sizes 3×3 and 7×7 , the exact box-filtered outputs are computed and compared with the box-filtered output estimates obtained using our framework, and RG-CoSamP as well. The heat map visualizations of the outputs are shown in figure 4. It is clear that greater quality box-filtered output estimates can be recovered using our framework and the recovered outputs retain the information regarding medium-sized structures in the images, while in the case of RG-CoSamP, the output is all noisy and does not give us any information.

5 TRACKING APPLICATION

In this section, we show the utility of the framework in practical applications. In particular, we show tracking results on 50 challenging videos used in benchmark comparison of tracking algorithms [37]. We emphasize that our aim is not to obtain state-of-the-art tracking results but to show that integral image estimates can be used to obtain robust tracking results at low measurement rates. To this end, we conduct two sets of tracking experiments, one with original resolution videos and one with high definition videos.

Tracking with original resolution videos: We conduct tracking experiments on original resolution videos at three different measurement rates, viz 1.28%, 1%, and 0.49%. In each case, we use the measurement matrix obtained for block size of 32×32 , and obtain **ReFINE** measurements for each frame using the ϕ^* obtained as above. Once, the measurements are obtained, our framework recovers integral image estimates from these measurements in real time. The estimates are subsequently fed into the best performing Haar-feature based tracking algorithm, Struck [38] to obtain the tracking results. Henceforth, we term this tracking pipeline as **ReFINE+Struck**. To evaluate our tracking results, we use the standard criterion of precision curve as suggested by [37]. To obtain the precision curve, we plot the precision scores against a range of thresholds. A frame contributes to the precision score for a particular threshold, α if the distance between the ground truth location of the target and estimated location by the tracker is less than the threshold, α . Precision score for given

Method	ReFlnE	RG-CoSamP	ReFlnE	RG-CoSamP	ReFlnE	RG-CoSamP
M (measurement ratio)	20 (0.005)	20 (0.005)	40 (0.01)	40 (0.01)	60 (0.015)	60 (0.015)
Time in s	0.0034	0.38	0.0036	0.58	0.0031	0.97
RSNR in dB	38.95	-16.76	38.96	-11.22	38.96	-10.9

TABLE 1

Comparison of average RSNR and time for recovered integral image estimates obtained using our method with RG-CoSamP. Our framework outperforms RG-CoSamP in terms of both recovery signal-to-noise ratio and time taken to estimate the integral image, at all measurement rates.



Fig. 1. (Four images (L-R: 'Two Men', 'Plane', 'Train' and 'Room') are randomly chosen for presenting qualitative and quantitative results.

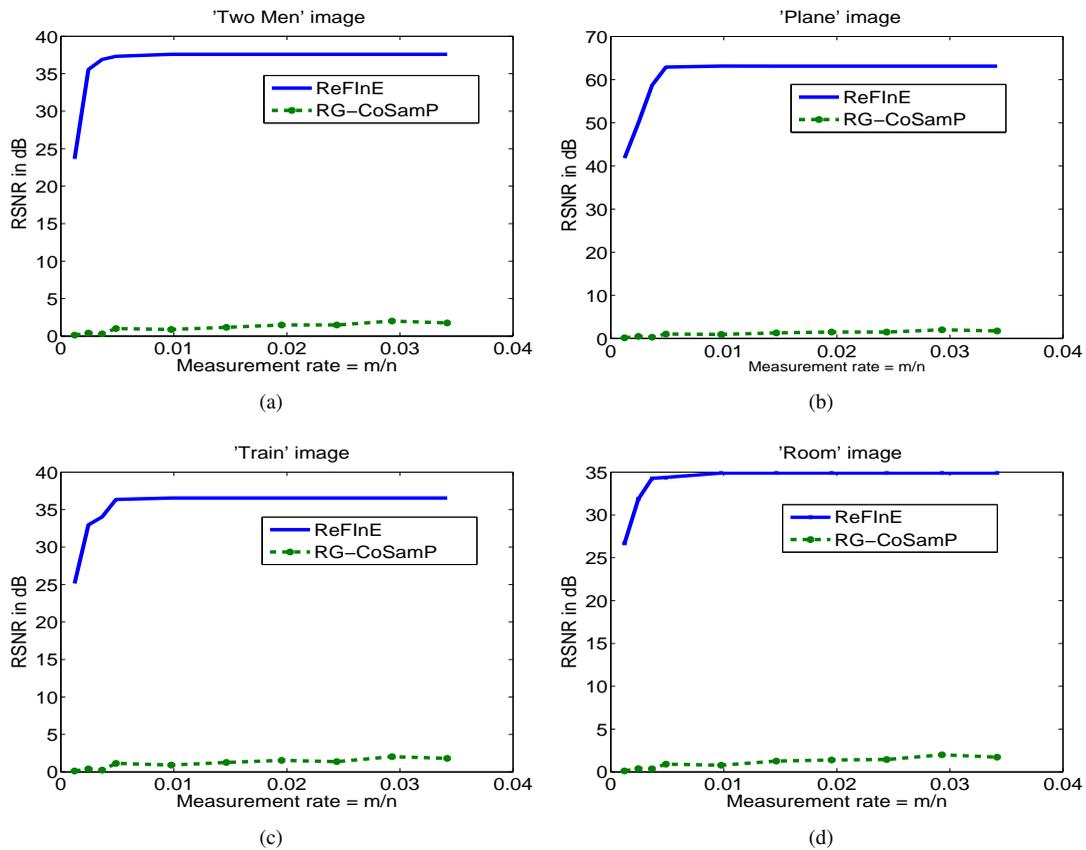


Fig. 2. The figure shows variation of image-wise RSNR for recovered integral image estimates for the four images. It is very clear that for all the four images, our framework outperforms 'RG-CoSamP' in terms of RSNR, at all measurement rates.

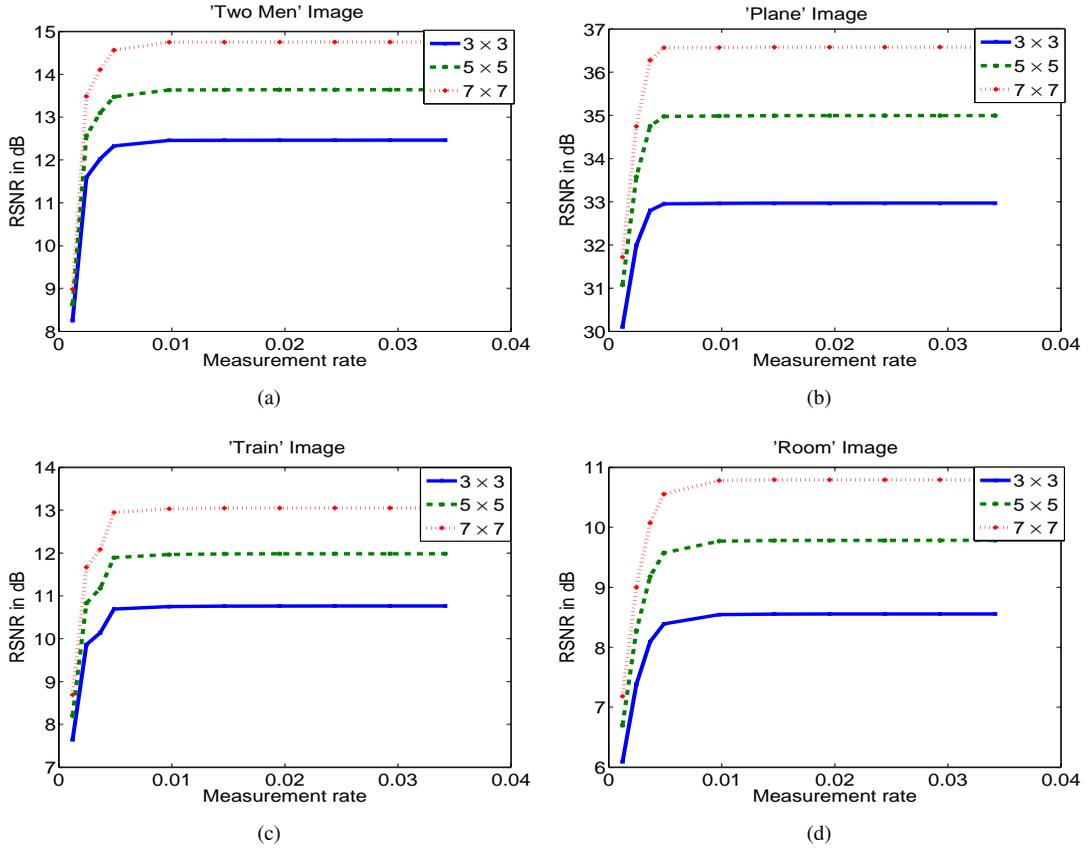


Fig. 3. The figure shows the variation of RSNR for the recovered box-filtered outputs using **ReFinE** with measurement rate. It is evident that even for 1% measurement rate, we obtain high RSNR box-filtered outputs. For a fixed measurement rate, the RSNR increases with the size of the filter. This shows the structures global in nature are captured better. This is particularly true in the case of 'Plane' image. The high RSNR for this image hints at the absence of fine structures and homogeneous background.

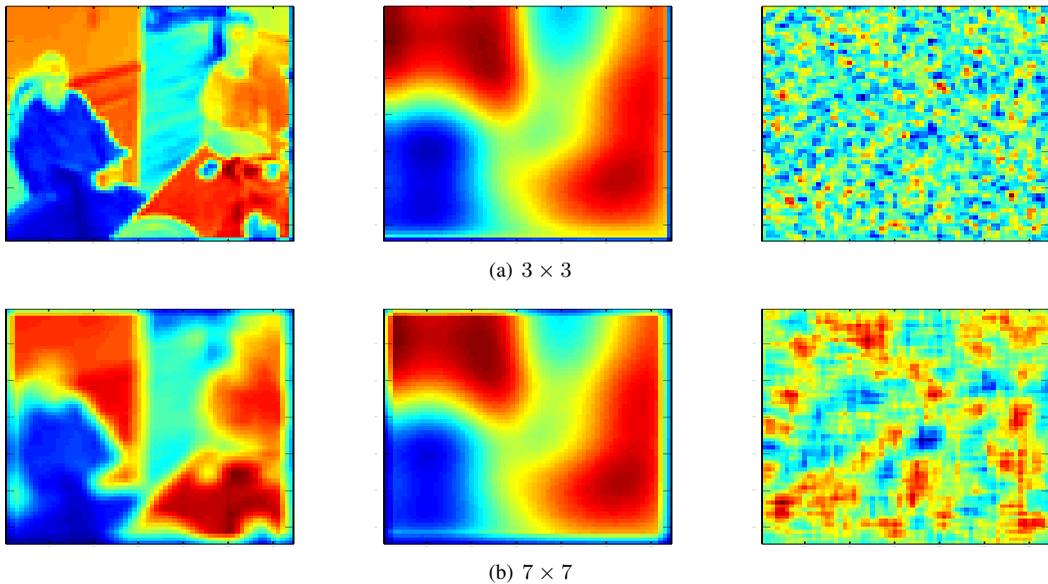


Fig. 4. Heat maps for box-filtered outputs for the 'Two men' image. Left to right: Exact output, **ReFinE** ($m/n = 0.01$), and RG-CosamP ($m/n = 0.01$). It is clear that greater quality box-filtered output estimates can be recovered from **ReFinE** measurements and the recovered outputs retain the information regarding medium-sized structures in the images, while in case of RG-CoSamP, the output is all noisy and does not give us any information.

threshold is calculated as the percentage of frames which contribute to the precision score. When precision scores are required to be compared with other trackers at one particular threshold, generally threshold is chosen to be equal to 20 pixels [37].

Precision curve: The precision curves for our framework at the three different measurement rates are plotted against a wide range of location error thresholds, and are compared with the same for other trackers, Oracle Struck [38], and various other trackers, TLD [20], MIL [19], CSK [39], and FCT [40] in figure 5. It is to be noted all the trackers used for comparison utilize full-blown images for tracking and hence operate at 100% measurement rate. As can be seen clearly, ‘ReFInE+Struck’ at 1.28% performs better than other trackers, MIL, CSK, TLD, and FCT and only a few percentage points worse than Oracle Struck for all thresholds. In particular, the mean precision over all 50 sequences in the dataset [37] for the threshold of 20 pixels is obtained for ‘ReFInE+Struck’ at three different measurement rates and is compared with other trackers in table 2. We obtain a precision of 59.26% at a measurement of 1.28%, which is only a few percentage points less than precision of 65.5% using Oracle Struck and 60.8% using TLD. Even at an extremely low measurement rate of 0.49%, we obtain mean precision of 45.78% which is competitive when compared to other trackers, MIL, and FCT which operate at 100% measurement rate. This clearly suggests that the small number of well-tailored measurements obtained using our framework retain enough information regarding integral images and hence also the Haar-like features which play a critical role in achieving tracking with high precision.

Frame rate: Even though, our framework uses Struck tracker, the frame rates at which ‘ReFInE+Struck’ operates are potentially less than the frame rate that can be obtained with Oracle Struck, and can even be different at different measurement rates. This is due to the fact that once the measurements are obtained for a particular frame, we first have to obtain an intermediate reconstructed frame before applying the integral operation. However, in the case of Oracle Struck, the integral operation is applied directly on the measured frame. The frame rate for ‘Our+Struck’ at different measurement rates are compared with the frame rates for other trackers in table 2. However, as can be seen, the preprocessing operation to obtain the intermediate reconstructed frame barely affects the speed of tracking since the preprocessing step, being multiplication of small-sized matrices can be efficiently at nearly 1000 frames per second.

Experiments with sequence attributes: Each video sequence in the benchmark dataset is annotated with a set of attributes, indicating the various challenges the video sequence offers in tracking. We plot precision percentages against the location error threshold for each of these 10 different kinds of attributes. Figure 6 shows the corresponding plots for attributes, ‘Illumination Variation’, ‘Background Clutter’, ‘Occlusion’, and ‘Scale Variation’. In the case of ‘Illumination Variation’ and ‘Occlusion’ ‘Our+Struck’ at measurement rate of 1.28% performs better

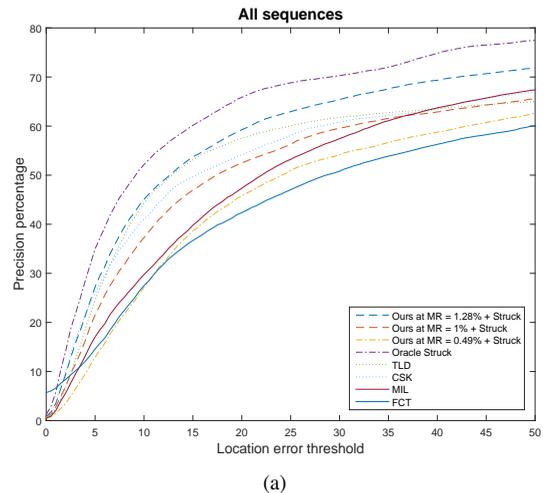


Fig. 5. ‘ReFInE+Struck’ at a measurement rate of 1.28% performs better than other trackers, MIL, CSK, TLD, and FCT and only a few percentage points worse than Oracle Struck for all thresholds. Even at a measurement rate of 1%, ‘ReFInE+Struck’ performs nearly as well as TLD and CSK trackers which operate at 100% measurement rate.

Tracker	Mean Precision	Mean FPS
ReFInE at MR = 1.28% + Struck	59.26	19.61
ReFInE at MR = 1% + Struck	52.47	19.61
ReFInE at MR = 0.49% + Struck	45.78	19.62
Oracle Struck [38]	65.5	20
TLD [20]	60.8	28
CSK [39]	54.11	362
MIL [19]	47.5	38
FCT [40]	42.37	34.92

TABLE 2

Mean precision percentage for 20 pixels error and mean frames per second for various state-of-the-art trackers are compared with our framework at different measurement rates. The precision percentages for our framework are stable even at extremely low measurement rates, and compare favorably with other trackers which operate at 100% measurement rate, i.e utilize all the pixels in the frames.

than TLD, CSK, MIL and FCT, whereas in the case of the ‘Background Clutter’ and ‘Scale Variation’ attributes, TLD performs slightly better than ‘Our+Struck’ at measurement rate of 1.28%.

Figure 7 shows the corresponding plots for attributes, ‘Deformation’, ‘Fast Motion’, ‘Motion Blur’, and ‘Low Resolution’. In the cases of ‘Deformation’, ‘Fast Motion’ and ‘Motion Blur’, ‘ReFInE+Struck’ at measurement rate of 1.28% performs better than TLD, CSK, MIL and FCT, whereas in the case of ‘Low Resolution’, TLD performs better than ‘ReFInE+Struck’.

Figure 8 shows the corresponding plots for attributes, ‘In the Plane rotation’, ‘Out of View’, and ‘Out of Plane rotation’.

Tracking with high resolution videos: Tracking using high-resolution videos can potentially lead to improvement in performance due to availability of fine-grained informa-

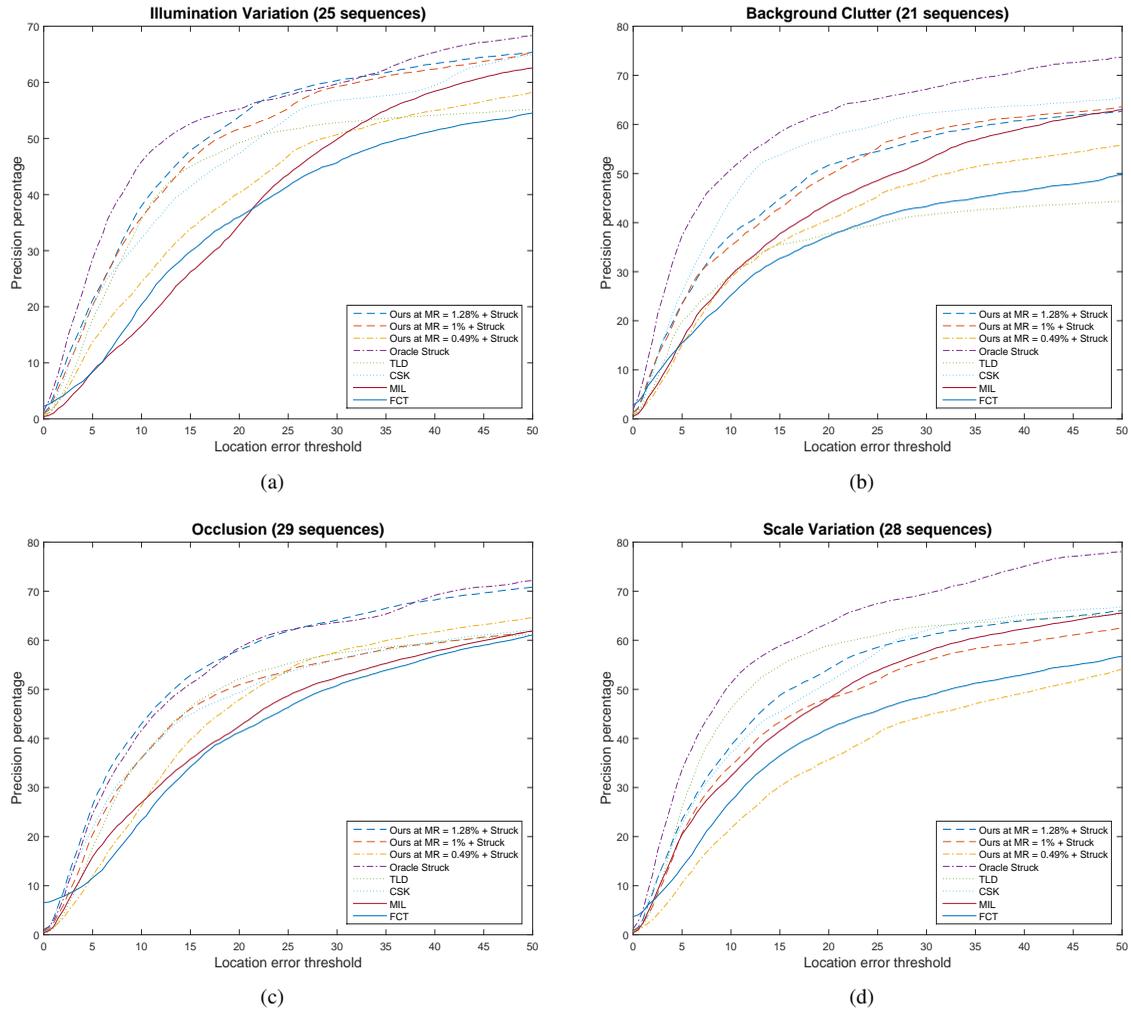


Fig. 6. Precision plots for four different attributes. In the case of ‘Illumination Variation’ and ‘Occlusion’ ‘ReFInE+Struck’ at measurement rate of 1.28% performs better than TLD, CSK, MIL and FCT, whereas in the case of the ‘Background Clutter’ and ‘Scale Variation’ attributes, TLD performs slightly better than ‘ReFInE+Struck’ at measurement rate of 1.28%.

tion regarding the scene. However, in many applications, the deployment of high-resolution sensors is severely limited by the lack of storage capacity. In such scenarios, it will be interesting to see if the small number of ReFInE measurements of high-resolution videos can yield better tracking performance than the full-blown low-resolution videos. To conduct tracking experiments on high resolution videos, we first employ a deep convolutional network based image super resolution (SR) algorithm, SRCNN, [41] to obtain high resolution frames of the 50 videos considered earlier in the section. The aspect ratio for all frames is maintained, and the upscaling factor for image super resolution is calculated such that the resolution of the longer dimension in the higher resolution frame is at least 1000 pixels. We found that upscale factors varies between 2 and 8 for various videos in the dataset. Once the high resolution videos are obtained, we proceed to obtain ReFInE measurements as before. We conduct tracking experiments at four different measurement rate (1%, 0.49%, 0.29%, 0.2%). Note that these different measurement rates are with respect to (wrt) the high-resolution frames, and

the measurement rate wrt original resolution, which we call effective measurement rate (EMR), is given by the ratio of the number of ReFInE measurements per frame to the number of pixels in a frame of original resolution video. Here, the tracking algorithm, Struck which we used for original resolution videos does not scale well in terms of computational complexity. For higher resolution videos, where the search space is much larger, we found that Struck is too slow for real-time application. Instead, we use a faster Haar feature based tracking algorithm, FCT [40] algorithm. Henceforth, we dub this tracking pipeline as **SR+ReFInE+FCT**. Once tracking results are obtained for the high resolution videos are obtained, we normalize the coordinates so as to obtain the tracking outputs with respect to original resolution videos. The precision score is calculated as before. The mean precision percentage for 20 pixels error and mean frames per second for ‘SR + ReFInE + FCT’ at various measurement rates are given in table 3 and are compared for the same for ‘Oracle FCT’ which operates for full-blown original resolution videos. It is clear that we obtain a significant boost in tracking accuracy

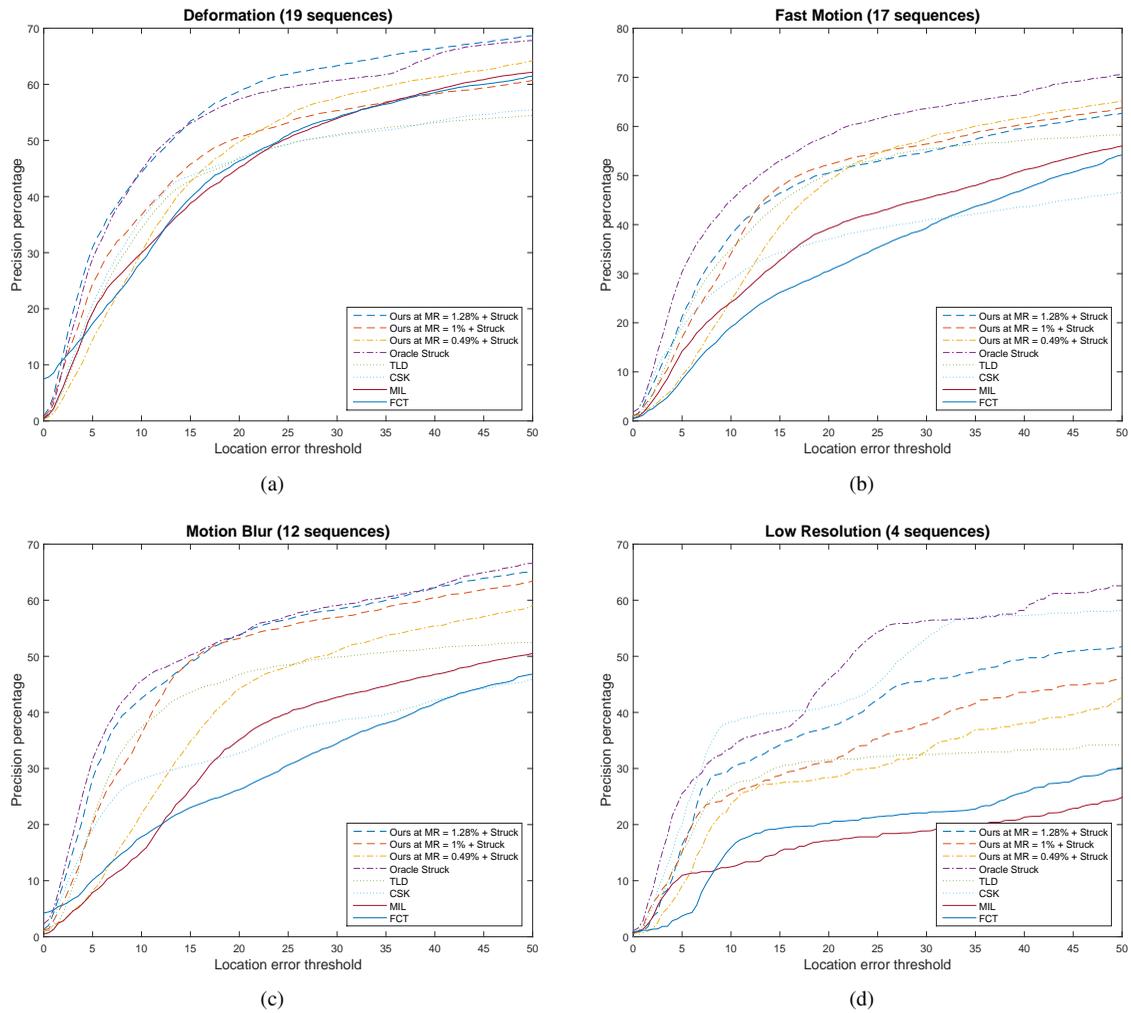


Fig. 7. Precision plots for four different attributes. In the cases of ‘Deformation’, ‘Fast Motion’ and ‘Motion Blur’, ‘ReFinE+Struck’ at measurement rate of 1.28% performs better than TLD, CSK, MIL and FCT, whereas in the case of ‘Low Resolution’, TLD performs better than ‘ReFinE+Struck’.

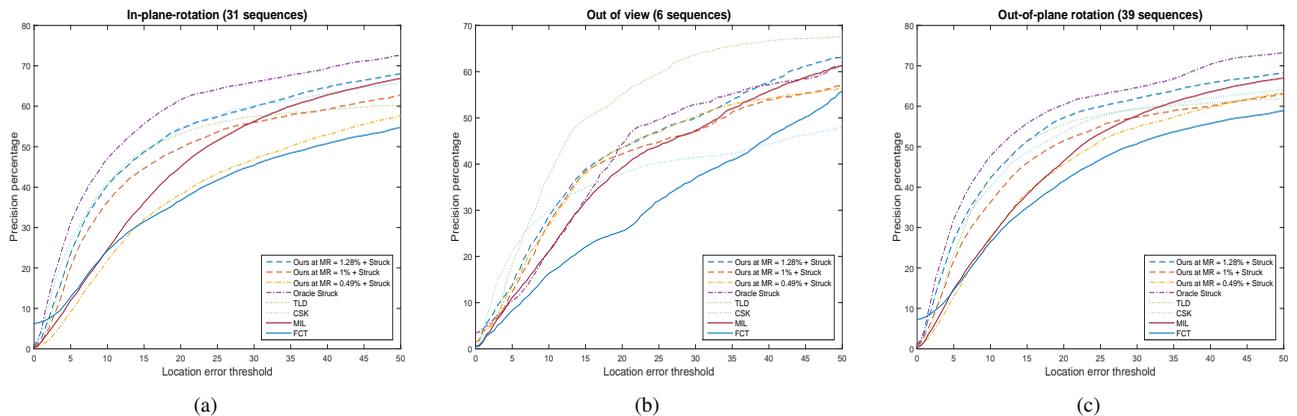


Fig. 8. Precision plots for three different attributes. In the cases of ‘In the plane rotation’, and ‘Out of plane rotation’, ‘ReFinE+Struck’ at measurement rate of 1.28% performs better than TLD, CSK, MIL and FCT, whereas in the case of ‘Out of View’, TLD performs better than ‘ReFinE+Struck’.

for high resolution videos. At measurement rate of 1% (EMR of 8.16%), we obtain a mean precision percentage of 54.83, which is 12.46 percentage points more than that for

‘Oracle FCT’. Even at a measurement rate of 0.2% (EMR of 1.63%), the precision percentage of 45.79, which is about 3.42 percentage points more than that for ‘Oracle

Tracker	Mean Precision	Mean FPS
SR + ReFlInE at EMR = 8.16% + FCT	54.83	19.61
SR + ReFlInE at EMR = 4% + FCT	53.03	19.61
SR + ReFlInE at EMR = 2.37% + FCT	50.9	19.62
SR + ReFlInE at EMR = 1.63% + FCT	45.79	19.62
Oracle FCT [40]	42.37	34.92

TABLE 3

Mean precision percentage for 20 pixels error and mean frames per second for ‘SR + ReFlInE + FCT’ at various measurement rates are compared with ‘Oracle FCT’. Even at extremely low measurement rates, the precision percentages for ‘SR + ReFlInE + FCT’ are better than for ‘Oracle FCT’ which operates on full-blown original resolution images.

FCT’. However, the more accurate precision comes at the cost of frame rate. Since the search space is much larger for high resolution videos, the speed of tracking for high resolution videos, is only about 20 FPS, while ‘Oracle FCT’ operates at 34.92 FPS. But 20 FPS suffices for near real-time implementations.

6 CONCLUSIONS

In this paper, we qualitatively and quantitatively showed that it is possible to obtain high quality estimates of integral images and box-filtered outputs directly from a small number of specially designed spatially multiplexed measurements called **ReFlInE** measurements. To show the practical applicability of the integral image estimates, we presented impressive reconstruction-free tracking results on challenging videos at an extremely low measurement rate of 1%. We also showed that with only a small number of **ReFlInE** measurements on high-resolution videos, which is only a fraction (2-8%) of the size of the original resolution, one can obtain significantly better object tracking results than using full blown original resolution videos. From a philosophical point of view, this points to the possibility of attaining greater performance on other computer vision inference tasks from a small number of carefully tailored spatially multiplexed measurements of high-resolution imagery rather than full-blown low resolution imagery.

APPENDIX A

DERIVATION: $\text{RANK}(\mathbf{P}) = \text{RANK}(\mathbf{Q}) - 1$

By construction we have $(\phi^d)^T = [\mathbf{1} | (\phi^d)_{2:m}^T] = [\mathbf{1} | (\mathbf{D}\mathbf{U}_{2:n}^T)^T]$. The column of all ones in $(\phi^d)^T$ is orthogonal to remaining $m - 1$ columns. Hence, we have $\text{rank}(\phi_{2:m}^d) = \text{rank}(\phi^d) - 1 = \text{rank}(\mathbf{Q}) - 1$. Similarly, we have $\text{rank}(\phi_{2:m}) = \text{rank}(\phi) - 1 = \text{rank}(\mathbf{Q}) - 1$. Since, \mathbf{P} is the product of equally ranked matrices, $(\phi^d)_{2:m}^T$ and ϕ , it follows that $\text{rank}(\mathbf{P}) = \text{rank}(\phi_{2:m}^d) = \text{rank}(\mathbf{Q}) - 1$.

APPENDIX B

CALCULATION OF $\mathcal{A}_i^*(\mathbf{y}_i^{k-1})$

For brevity, we drop the superscript, $k - 1$. Let $\Sigma_{w_d}^{1/2} \mathbf{U}_{2:n}^T = \Sigma_U$. Consider the following equation.

$$\langle \mathcal{A}(\mathbf{P}), \mathbf{y} \rangle = \sum_{i=1}^n \langle \mathbf{P}, \mathcal{A}_i^*(\mathbf{y}_i) \rangle. \quad (18)$$

The left hand side can be written as

$$[\mathbf{y}_1, \dots, \mathbf{y}_n] [(\mathcal{A}_1(\mathbf{P}))^T, \dots, (\mathcal{A}_n(\mathbf{P}))^T]^T \quad (19)$$

$$= [\mathbf{y}_1, \dots, \mathbf{y}_n] [h_1^T \mathbf{P}^T \Sigma_U^T, \dots, h_n^T \mathbf{P}^T \Sigma_U^T]^T \quad (20)$$

$$= \sum_{i=1}^n \mathbf{y}_i \Sigma_U \mathbf{P}^T h_i \quad (21)$$

$$= \sum_{i=1}^n \langle \mathbf{P}, \Sigma_U^T \mathbf{y}_i h_i^T \rangle. \quad (22)$$

Comparing the equations 18 and 22, we have $\mathcal{A}_i^*(\mathbf{y}_i) = \Sigma_U^T \mathbf{y}_i h_i^T$.

ACKNOWLEDGMENTS

REFERENCES

- [1] S. K. Nayar, V. Branzoi, and T. E. Boult, “Programmable imaging: Towards a flexible camera,” *Intl. J. Comp. Vision*, vol. 70, no. 1, pp. 7–22, 2006.
- [2] M.B. Wakin, J.N. Laska, M.F. Duarte, D. Baron, S. Sarvotham, D. Takhar, K.F. Kelly and R.G. Baraniuk, “An architecture for compressive imaging,” in *IEEE Conf. Image Process.*, 2006.
- [3] A. C. Sankaranarayanan, C. Studer, and R. G. Baraniuk, “Cs-muvi: Video compressive sensing for spatial-multiplexing cameras,” in *Computational Photography (ICCP), 2012 IEEE International Conference on.* IEEE, 2012, pp. 1–10.
- [4] A. C. Sankaranarayanan, P. Turaga, R. Baraniuk, and R. Chellappa, “Compressive acquisition of dynamic scenes,” in *Euro. Conf. Comp. Vision*, 2010.
- [5] D. L. Donoho, “Compressed sensing,” *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [6] E. J. Candes and T. Tao, “Near-optimal signal recovery from random projections: Universal encoding strategies?” *IEEE Trans. Inf. Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.
- [7] J. A. Tropp and A. C. Gilbert, “Signal recovery from random measurements via orthogonal matching pursuit,” *IEEE Trans. Inf. Theory*, vol. 53, no. 12, pp. 4655–4666, 2007.
- [8] D. Needell and J. A. Tropp, “Cosamp: Iterative signal recovery from incomplete and inaccurate samples,” *Applied and Computational Harmonic Analysis*, vol. 26, no. 3, pp. 301–321, 2009.
- [9] V. Thirumalai and P. Frossard, “Correlation estimation from compressed images,” *J. Visual Communication and Image Representation*, vol. 24, no. 6, pp. 649–660, 2013.
- [10] K. Kulkarni and P. Turaga, “Recurrence textures for activity recognition using compressive cameras,” in *IEEE Conf. Image Process.*, 2012.
- [11] M. A. Davenport, M. F. Duarte, M. B. Wakin, J. N. Laska, D. Takhar, K. F. Kelly and R. G. Baraniuk, “The smashed filter for compressive classification and target recognition,” pp. 6498–6499, 2007.
- [12] R. Calderbank, S. Jafarpour and R. Schapire, “Compressed learning: Universal sparse dimensionality reduction and learning in the measurement domain,” in *Preprint*, 2009.
- [13] K. Kulkarni and P. Turaga, “Reconstruction-free action inference from compressive imagers,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PP, no. 99, 2015.
- [14] S. Lohit, K. Kulkarni, P. Turaga, J. Wang, and A. Sankaranarayanan, “Reconstruction-free inference on compressive measurements,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2015, pp. 16–24.
- [15] P. Viola and M. J. Jones, “Robust real-time face detection,” *Intl. J. Comp. Vision*, vol. 57, no. 2, pp. 137–154, 2004.

- [16] P. Dollár, S. Belongie, and P. Perona, "The fastest pedestrian detector in the west." in *British Machine Vision Conf.*, vol. 2, no. 3, 2010, p. 7.
- [17] H. Grabner, M. Grabner, and H. Bischof, "Real-time tracking via on-line boosting." in *British Machine Vision Conf.*, 2006.
- [18] K. Zhang, L. Zhang, and M.-H. Yang, "Real-time compressive tracking," in *Euro. Conf. Comp. Vision*, 2012, pp. 864–877.
- [19] B. Babenko, M.-H. Yang, and S. Belongie, "Robust object tracking with online multiple instance learning," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 8, pp. 1619–1632, 2011.
- [20] Z. Kalal, J. Matas, and K. Mikolajczyk, "Pn learning: Bootstrapping binary classifiers by structural constraints," in *IEEE Conf. Comp. Vision and Pattern Recog.*, 2010, pp. 49–56.
- [21] S. Ramakanth and R. Babu, "Seamseg: Video object segmentation using patch seams," in *IEEE Conf. Comp. Vision and Pattern Recog.*, 2013, pp. 376–383.
- [22] J. M. Duarte-Carvajalino and G. Sapiro, "Learning to sense sparse signals: Simultaneous sensing matrix and sparsifying dictionary optimization," *IEEE Transactions on Image Processing*, vol. 18, no. 7, pp. 1395–1408, 2009.
- [23] T. Goldstein, L. Xu, K. F. Kelly, and R. Baraniuk, "The stone transform: Multi-resolution image enhancement and real-time compressive video," *arXiv preprint arXiv:1311.3405*, 2013.
- [24] H. S. Chang, Y. Weiss, and W. T. Freeman, "Informative sensing," *arXiv preprint arXiv:0901.4275*, 2009.
- [25] M. A. Davenport, M. F. Duarte, M. B. Wakin, J. N. Laska, D. Takhar, K. F. Kelly and R. G. Baraniuk, "The smashed filter for compressive classification and target recognition," *Computat. Imag. V.*, vol. 6498, pp. 142–153, 2007.
- [26] D. F. Andrews and C. L. Mallows, "Scale mixtures of normal distributions," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 99–102, 1974.
- [27] M. J. Wainwright and E. P. Simoncelli, "Scale mixtures of gaussians and the statistics of natural images." in *Adv. Neural Inf. Proc. Sys.*, 1999, pp. 855–861.
- [28] S. Lyu and E. P. Simoncelli, "Modeling multiscale subbands of photographic images with fields of gaussian scale mixtures," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 4, pp. 693–706, 2009.
- [29] S. G. Mallat, "A theory for multiresolution signal decomposition: the wavelet representation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 11, no. 7, pp. 674–693, 1989.
- [30] G. Frahm, "Generalized elliptical distributions: theory and applications," Ph.D. dissertation, Universität zu Köln, 2004.
- [31] E. J. Candes, T. Strohmer, and V. Voroninski, "Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming," *Communications on Pure and Applied Mathematics*, vol. 66, no. 8, pp. 1241–1274, 2013.
- [32] C. Hegde, A. C. Sankaranarayanan, W. Yin, and R. G. Baraniuk, "Numax: A convex approach for learning near-isometric linear embeddings."
- [33] Y.-J. Liu, D. Sun, and K.-C. Toh, "An implementable proximal point algorithmic framework for nuclear norm minimization," *Mathematical programming*, vol. 133, no. 1-2, pp. 399–436, 2012.
- [34] J.-F. Cai, E. J. Candès, and Z. Shen, "A singular value thresholding algorithm for matrix completion," *SIAM Journal on Optimization*, vol. 20, no. 4, pp. 1956–1982, 2010.
- [35] M. Fukushima, Z.-Q. Luo, and P. Tseng, "Smoothing functions for second-order-cone complementarity problems," *SIAM Journal on optimization*, vol. 12, no. 2, pp. 436–460, 2002.
- [36] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results," <http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html>.
- [37] Y. Wu, J. Lim, and M.-H. Yang, "Online object tracking: A benchmark," in *IEEE Conf. Comp. Vision and Pattern Recog.* IEEE, 2013, pp. 2411–2418.
- [38] S. Hare, A. Saffari, and P. H. Torr, "Struck: Structured output tracking with kernels," in *IEEE Intl. Conf. Comp. Vision.* IEEE, 2011, pp. 263–270.
- [39] J. F. Henriques, R. Caseiro, P. Martins, and J. Batista, "Exploiting the circulant structure of tracking-by-detection with kernels," in *Euro. Conf. Comp. Vision*, 2012, pp. 702–715.
- [40] K. Zhang, L. Zhang, and M. Yang, "Fast compressive tracking," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 36, no. 10, pp. 2002–2015, 2014.
- [41] C. Dong, C. C. Loy, K. He, and X. Tang, "Learning a deep convolutional network for image super-resolution," in *Euro. Conf. Comp. Vision.* Springer, 2014, pp. 184–199.



Kuldeep Kulkarni



Pavan Turaga