INSENSE: INCOHERENT SENSOR SELECTION FOR SPARSE SIGNALS

Amirali Aghazadeh[1], *Mohammad Golbabaee*[2], *Andrew Lan*[3], *Richard Baraniuk*[4]

[1]Stanford University, [2]University of Edinburgh, [3]Princeton University, [4]Rice University

ABSTRACT

Sensor selection refers to the problem of intelligently selecting a small subset of a collection of available sensors to reduce the sensing cost while preserving signal acquisition performance. The majority of sensor selection algorithms find the subset of sensors that best recovers an arbitrary signal from a number of linear measurements that is larger than the dimension of the signal. In this paper, we develop a new sensor selection algorithm for *sparse* (or near sparse) signals that finds a subset of sensors that best recovers such signals from a number of measurements that is much smaller than the dimension of the signal. Existing sensor selection algorithms cannot be applied in such situations. Our proposed *Incoherent Sensor Selection (Insense)* algorithm minimizes a coherence-based cost function that is adapted from recent results in sparse recovery theory. Using three datasets, including a real-world dataset on microbial diagnostics, we demonstrate the superior performance of Insense for sparse-signal sensor selection.

Index Terms— Sensor selection, coherence, optimization, compressive sensing

1. INTRODUCTION

The accelerating demand for capturing signals at high resolution is driving acquisition systems to employ an increasingly large number of sensing units. However, factors like manufacturing costs, physical limitations, and energy constraints typically define a budget on the total number of sensors that can be implemented in a given system. This budget constraint motivates the design of *sensor selection* algorithms [1] that intelligently select a subset of sensors from a pool of available sensors in order to lower the sensing cost with only a small deterioration in acquisition performance.

In this paper, we extend the classical sensor selection setup, where D available sensors obtain linear measurements of a signal $x \in \mathbb{R}^N$ according to $y = \Phi x$ with each row of $\Phi \in \mathbb{R}^{D\times N}$ corresponding to one sensor. In this setup, the sensor selection problem is one of finding a subset Ω of sensors (i.e., rows of Φ) of size $|\Omega| = M$ such that the signal x can be recovered from its M linear measurements

$$y_\Omega = \Phi_\Omega x \tag{1}$$

with minimal reconstruction error. Here, $\Phi_\Omega \in \mathbb{R}^{M\times N}$ is called the *sensing matrix*; it contains the rows of Φ indexed by Ω. The lion's share of current sensor selection algorithms [1–3] select sensors that best recover an arbitrary signal x from $M > N$ measurements. In this case, (1) is *overdetermined*. Given a subset of sensors Ω, the signal x is recovered simply by inverting the sensing matrix while computing $\Phi_\Omega^\dagger y_\Omega$, where $\Phi_\Omega^\dagger$ is the pseudoinverse of Φ_Ω. Such approaches do not exploit the fact that many real-world signals are (near) *sparse* in some basis [4]. It is now well-known that (near) sparse signals can be accurately recovered from a number of linear measurements $M \ll N$ using sparse recovery/compressive sensing (CS) techniques [5–7]. Conventional sensor selection algorithms are not designed to exploit low-dimensional signal structure. Indeed, they typically fail to select the appropriate sensors for sparse signals in this *underdetermined* setting ($M < N$).

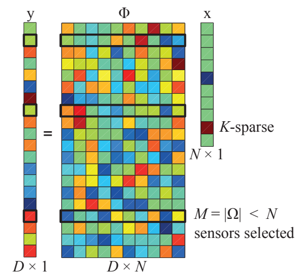

Fig. 1: Schematic of the sensor selection problem for sparse signals. Here, $M = 3$ sensors indexed by $\Omega = \{2, 8, 17\}$ are selected from $D = 20$ available sensors to recover a $K = 2$-sparse vector $x \in \mathbb{R}^{N=10}$, from the linear system $y_\Omega = \Phi_\Omega x$.

In this paper, we develop a new sensor selection framework that finds the optimal subset of sensors Ω that best recovers a (near) sparse signal x from $M < N$ linear measurements (see Fig. 1). In contrast to the conventional sensor selection setting, here the sensing equation (1) is underdetermined, and it can not be simply inverted in closed form. A key challenge in sensor selection in the underdetermined setting is that we must replace the cost function that has been useful in the classical, overdetermined setting, namely the estimation error $\|x - \hat{x}\|_2^2$ (or the covariance of the estimation error in the presence of noise). In the overdetermined setting, this error can be obtained in closed form simply by inverting (1). In the underdetermined setting, this error has no closed form expression. Indeed, recovery of a sparse vector x from y_Ω requires a computational scheme [8, 9].

978-1-5386-4658-8/18/\$31.00 ©2018 IEEE

Fortunately, the sparse recovery theory tells us that one can reliably recover a sufficiently sparse vector x from its linear measurements y_Ω when the columns of the sensing matrix Φ_Ω are sufficiently *incoherent* [10–12]. Define the *coherence* between the columns ϕ_i and ϕ_j in the sensing matrix Φ_Ω as $\mu_{ij}(\Phi_\Omega) = \frac{|\langle \phi_i, \phi_j \rangle|}{\|\phi_i\|\|\phi_j\|}$. If the values of $\mu_{ij}(\Phi_\Omega)$ for all pairs of columns (i, j) are bounded by a certain threshold, then sparse recovery algorithms such as Basis Pursuit (BP) [10, 13, 14] can recover the sparse signal x exactly. This theory suggests a new cost function for sensor selection. To select the sensors Ω that reliably recover a sparse vector, we can minimize any monotonic function of coherence. Here we minimize the *average squared coherence* $\mu^2_{\text{avg}}(\Phi_\Omega) = \frac{1}{\binom{N}{2}} \sum_{1 \le i < j \le N} \mu^2_{ij}(\Phi_\Omega)$, to maximize the *average* recovery performance over sparse vectors x. The challenge now becomes formulating an optimization algorithm that selects the subset of the *rows* of Φ (the sensors) whose *columns* have the smallest average squared coherence.

1.1. Contributions

We make three distinct contributions in this work. First, we propose the sparse-signal sensor selection problem and we demonstrate that the standard cost functions used in overdetermined sensor selection algorithms are not suitable for the underdetermined case.

Second, we develop a new sensor selection algorithm that optimizes the new cost function $\mu^2_{\text{avg}}(\Phi_\Omega)$; call it the *Incoherent Sensor Selection (Insense)* algorithm. Insense employs an efficient optimization technique to find a subset of sensors with smallest average coherence among the columns of the selected sensing matrix Φ_Ω. The optimization technique – projection onto the convex set defined by a scaled-boxed simplex (SBS) constraint – is of independent interest. We have made the codes for the Insense algorithm available online at https://github.com/amirmohan/Insense.git.

Third, we demonstrate the superior performance of Insense over conventional sensor selection algorithms using an exhaustive set of experimental evaluations that include real-world datasets from microbial diagnostics and four performance metrics: average mutual coherence, sparse recovery performance, frame potential, and condition number. We demonstrate that, for the kinds of redundant, coherent, or structured Φ that are common in real-world applications, Insense finds the best subset of sensors in terms of sparse recovery performance. Indeed, in these cases, many conventional sensor selection algorithms fail completely.

1.2. Related work

Existing sensor selection algorithms mainly study the sensor selection problem in the overdetermined regime (when $M \ge N$) [1–3, 15]. In the overdetermined regime, robust signal recovery can be obtained using the solution to the least squares (LS) problem in the sensing model (1), which motivates as a cost function the mean squared error (MSE) [16–18] or a proxy of the MSE [19–21] of the LS solution. For instance, Joshi, et al. [1] employ a convex optimization-based algorithm to minimize the log-volume of the *confidence ellipsoid* around the LS solution of x. Shamaiah et al. [2] develop a greedy algorithm that outperforms the convex approach in terms of MSE. FrameSense [3] minimize the frame potential (FP) of the selected matrix $\text{FP}(\Phi_\Omega) = \sum_{\forall (i,j) \in \Omega,\, i<j} |\langle \phi^i, \phi^j \rangle|^2$, where ϕ^i represents the i^{th} row of Φ. Several additional sensor selection algorithms that assume a non-linear observation model [22, 23] also operate only in the overdetermined regime.

2. PROBLEM STATEMENT

Consider a set of D sensors taking nonadaptive, linear measurements of a K-sparse (i.e., with K non-zero elements) vector $x \in \mathbb{R}^N$ following the linear system $y = \Phi x$, where $\Phi \in \mathbb{R}^{D \times N}$ is a given sensing matrix. We aim to select a subset Ω of sensors of size $|\Omega| = M \ll D$, such that the sparse vector x can be recovered from the resulting $M < N$ linear measurements $y_\Omega = \Phi_\Omega x$ with minimal reconstruction error. Here, Φ_Ω contains the rows of Φ indexed by Ω, and y_Ω contains the entries of y indexed by Ω. This model for the sensor selection problem can be adapted to more general scenarios. For example, if the signal is sparse in a basis Ψ, then we simply consider $\Phi = \Theta\Psi$ as the new sensing matrix, where Θ is the original sensing matrix. In order to find a subset Ω of sensors (rows of Φ) that best recovers a sparse signal x from y_Ω,[1] we aim to select a submatrix $\Phi_\Omega \in \mathbb{R}^{M \times N}$ that attains the lowest average squared coherence $\mu^2_{\text{avg}}(\Phi_\Omega)$. The term μ_{avg} averages the off-diagonal entries of $\Phi_\Omega^T \Phi_\Omega$ (indexed by $1 \le i < j \le N$) after column normalization. Define the diagonal selector matrix $Z = \text{diag}(z)$ with $z = [z_1, z_2, z_3, \ldots, z_D]^T$ and $z_i \in \{0, 1\}$, where $z_i = 1$ indicates that the i^{th} row (sensor) in Φ is selected and $z_i = 0$ otherwise. This enables us to formulate the sensor selection problem as the following optimization problem $\min_{z \in \{0,1\}^D} \sum_{1 \le i < j \le N} \frac{{G_{ij}}^2}{G_{ii}\, G_{jj}}$, subject to $G = \Phi^T Z \Phi$, $\mathbf{1}^T z = M$, where $\mathbf{1}$ is the all-ones vector. This Boolean optimization problem is combinatorial, since one needs to search over $\binom{D}{M}$ combinations of sensors to find the optimal set Ω. To overcome this complexity, we relax the Boolean constraint on z_i to the *box constraint* $z_i \in [0, 1]$ to arrive at the following problem

$$\min_{z \in [0,1]^D} \sum_{1 \le i < j \le N} \frac{{G_{ij}}^2}{G_{ii}\, G_{jj}}, \text{s.t. } G = \Phi^T Z \Phi,\ \mathbf{1}^T z = M, \quad (2)$$

which supports an efficient gradient–projection algorithm to find an approximate solution. We develop this algorithm next.

[1] Or find one of the solutions if many solutions exists.

3. THE INSENSE ALGORITHM

We now outline the steps that Insense takes to solve the problem (2). We slightly modify the objective of (2) to

$$f_\epsilon(z) = \sum_{1\leq i<j\leq N} \frac{{G_{ij}}^2 + \epsilon_1}{G_{ii}\; G_{jj} + \epsilon_2} \text{ where } G = \Phi^T Z \Phi, \quad (3)$$

where the small positive constants $\epsilon_2 < \epsilon_1 \ll 1$ make the objective well-defined and bounded over $z \in [0,1]^D$. The objective function (3) is smooth and differentiable but non-convex; the box constraints on z are linear. We minimize the objective using the iterative gradient-projection algorithm. The gradient $\nabla_z f \in \mathbb{R}^D$ can be computed using the multidimensional chain rule of derivatives [24] as $(\nabla_z f)_i = \left(\Phi \nabla_G f \Phi^T\right)_{ii}$ for $i = 1, \ldots, D$. The $N \times N$ upper triangular matrix $\nabla_G f$ is the gradient of f in terms of the (auxiliary) variable G at the point $G = \Phi^T Z \Phi$, given by $(\nabla_G f)_{ij} = 2G_{ij}/(G_{ii}\; G_{jj} + \epsilon_2)$, when $i < j$, $(\nabla_G f)_{ij} = -\sum_{\forall \ell \neq i} G_{\ell\ell}\left(G_{i\ell}^2 + \epsilon_1\right)/(G_{ii}\; G_{\ell\ell} + \epsilon_2)^2$, when $i = j$, and $(\nabla_G f)_{ij} = 0$ elsewhere.

The Insense algorithm proceeds as follows. First, the variables G and Z are initialized randomly. Next, we perform the following update in iteration k, $z^{k+1} = P_{\text{SBS}}\left(z^k - \gamma^k \nabla_z f(z^k)\right)$, where P_{SBS} denotes the projection onto the convex set defined by the scaled boxed-simplex (SBS) constraints $\mathbf{1}^T z = M$ and $z \in [0,1]^D$ which we detail in the next paragraph. For certain bounded step size rules (e.g., $\gamma^k \leq 1/L$, where L is the Lipschitz constant of $\nabla_z f$), the sequence $\{z^k\}$ converges to a critical point of the non-convex problem [25, 26]. In our implementation, we use a backtracking line search to determine γ^k in each step [26]. We now detail our approach to solving the SBS projection problem

$$\min_z \frac{1}{2}\|z - y\|_2^2, \text{s.t.} \sum_i z_i = M, z_i \in [0,1]_{\forall i=1,\ldots,D}. \quad (4)$$

We emphasize that, for $M > 1$, the SBS projection problem is significantly different from the (scaled-)simplex constraint ($\sum_i z_i = M$) projection problem that has been studied in the literature [27–29], due to the additional box constraints $z_i \in [0,1]$. The Lagrangian of the problem (4) can be written as $f(z, \lambda, \alpha, \beta) = \frac{1}{2}\|z - y\|_2^2 \quad +\lambda\left(\sum_i z_i - M\right) + \sum_i \alpha_i(-z_i) + \sum_i \beta_i(z_i - 1)$, where λ, α, β are Lagrange multipliers for the equality and inequality constraints, respectively. The Karush-Kuhn-Tucker (KKT) conditions are given by $z_i - y_i + \lambda - \alpha_i + \beta_i = 0, \forall i$, $\sum_i z_i - M = 0$, $\alpha_i(-z_i) = 0$, $\beta_i(z_i - 1) = 0$, $\alpha_i, \beta_i \geq 0$, $0 \leq z_i \leq 1$, $\forall i$. According to the complimentary slackness condition for the box constraint $z_i \in [0,1]$, we have the following three cases for z_i: (a) $z_i = 0$: $\beta_i = 0, \alpha_i = y_i + \lambda > 0$, (b) $z_i = 1$: $\alpha_i = 0, \beta_i = 1 - y_i - \lambda > 0$, (c) $z_i \in (0,1)$: $\alpha_i = \beta_i = 0, z_i = y_i + \lambda$. Therefore, the value of λ holds the key to the proximal problem (4). However, finding λ is not an easy task, since we do not know which entries of z will fall on the boundary of the box constraint (and are equal to either 0 or 1). In order to find the entries z_i that are equal to 0, we assume without loss of generality that the values of y are sorted in ascending order: $y_1 \leq y_2 \leq \ldots y_D$. We note that, in all three cases above, $z_i = \max(\min(y_i - \lambda, 1), 0)$. Therefore, $\sum_i z_i$ is a non-decreasing function of λ. We evaluate $\sum_i z_i$ at the following values of λ: $\lambda = -y_1$: $z_1 = 0, z_i = \min(y_i - y_1, 1)$ for $i \geq 2$, $\lambda = -y_2$: $z_1 = z_2 = 0$, $z_i = \min(y_i - y_1, 1)$ for $i \geq 3$, $\ldots$ $\lambda = -y_D$: $z_1 = z_2 = \ldots = z_D = 0$. Thus, the entries in z that are equal to 0 correspond to the first K_0 entries in y, where K_0 is the largest integer k such that $\sum_i \max(\min(y_i - y_k, 1), 0) \geq M$. Similarly, we can find the entries in z that are equal to 1 by negating z and y in (4). Let $p = -y$ and assume that its entries are sorted in ascending order; a procedure similar to that above shows that the entries in z that are equal to 1 correspond to the first K_1 entries in p, where K_1 is the largest integer K such that $\sum_i \max(\min(p_i - p_k - 1, 0), -1) \geq -M$. Knowing which entries in z are equal to 0 and 1, we can solve for the value of λ by working with the entries with values in $(0,1)$. Using case (c) above and denoting the index set of these entries by ζ, we have $\lambda = (M - K_1 - \sum_{i\in\zeta} y_i)/|\zeta|$, and the solution to (4) is given by $z_i = \max(\min(y_i - \lambda, 1), 0)$.

Table 1: Comparison of Insense against the baseline algorithms on selecting $M = 10$ rows from a structured Identity/Gaussian Φ.

Algorithms	$\mu_{\text{avg}(\Phi_\Omega)}$	FP(Φ_Ω)	CN(Φ_Ω)	BP accuracy %
Insense	**0.3061±0.0047**	1019 ±313	1.93±0.19	**92.27±1.42**
EigenMaps	–	**0.00±0.00**	**1.00±0.00**	4.00±0.00
MSE-G	0.3872±0.0305	1155±374	11.51±0.93	57.91±1.09
FrameSense	–	**0.00±0.00**	**1.00±0.00**	4.00±0.00
MI-G	–	**0.00±0.00**	**1.00±0.00**	4.00±0.00
Entropy-G	–	**0.00±0.00**	**1.00±0.00**	4.00±0.00
Determinant-G	–	**0.00±0.00**	**1.00±0.00**	4.00±0.00
Greedy SS	–	**0.00±0.00**	**1.00±0.00**	4.00±0.00
Convex SS	0.3137±0.0075	2279±470	2.22±0.25	88.64±3.64

4. EXPERIMENTS

In this section, we experimentally validate Insense using a range of synthetic and real-world datasets. In all experiments, we set $\epsilon_1 = 10^{-9}$ and $\epsilon_2 = 10^{-10}$ (anything in the range $\epsilon_2 < \epsilon_1 \ll 1$ can be utilized). We terminate Insense when the relative change of the cost function $\mu^2_{\text{avg}}(\Phi_\Omega)$ drops below 10^{-7}. We compare Insense with several leading sensor selection algorithms, including *Convex Sensor Selection* [1], *Greedy Sensor Selection* [2], *EigenMaps* [15], and *FrameSense* [3]. We also compare with four greedy sensor selection algorithms that were featured in [3]. The first three minimize different information theoretic measures of the selected sensing matrix as a proxy to the MSE: the determinant in *Determinant-G* [19], mutual information (MI) in *MI-G* [20], and entropy in *Entropy-G* [21]. The final greedy algorithm, *MSE-G* [16–18], directly minimizes the MSE of the LS reconstruction error.

We compare the sensor selection algorithms using the following four metrics: *Average coherence* $\mu_{\text{avg}}(\Phi_\Omega)$, *Frame po-*

Table 2: Comparison of Insense against the baseline algorithms on selecting $M = 10$ rows from a structured Uniform/Gaussian Φ.

Algorithms	$\mu_{\text{avg}}(\Phi_\Omega)$	FP(Φ_Ω)	CN(Φ_Ω)	Gaussian %	BP accuracy %
Insense	**0.3165±0.0023**	9320 ±3292	**1.46±0.07**	**100±0**	**58.55±2.64**
EigenMaps	0.3215±0.0021	7230±2319	2.07±0.12	90±0	57.60±3.72
MSE-G	0.5805±0.0440	78530±12450	5.99±0.31	17±4	49.90±3.54
FrameSense	0.3273±0.0059	**6095±1708**	3.19±0.92	84±5	58.15±2.26
MI-G	0.6814±0.0556	93260±109250	6.26±0.77	7±4	51.60±5.21
Entropy-G	0.7007±0.0804	98950±16216	6.61±0.48	5±7	53.70±5.21
Determinant-G	0.7303±0.0545	105700±11228	6.57±0.31	3±4	55.50±4.50
Greedy SS	0.7303±0.0545	105700±11228	5.57±0.31	3±4	55.50±4.50
Convex SS	0.5788±0.1140	75270±27383	5.97±0.77	20±15	54.40±4.20

tential FP(Φ_Ω), *Condition number* CN(Φ_Ω), and *BP recovery accuracy*. Depending on the task, in some experiments we only report a subset of the metrics. To compute BP recovery accuracy, we average the performance of the BP algorithm [30] over 100 random sparse recovery trials.

Dataset 1) Identity/Gaussian matrix: We construct our first dataset by concatenating two 50×50 matrices: An identity matrix and a random matrix with i.i.d. Gaussian entries. Such matrices feature prominently in certain real-world CS problems (see [31]). To achieve an optimal sparse recovery performance, the sensor selection algorithm should select rows (sensors) from the Gaussian submatrix. Table 1 compares the performance of Insense to the baseline algorithms for the problem of selecting $M = 10$ rows from the structured Identity/Gaussian Φ. We repeat the same experiment 10 times with different random Gaussian matrices.[2] In particular, Insense, Convex SS, and MSE-G are the only algorithms that select rows of the Gaussian sub-matrix. While achieving the minimum FP(Φ_Ω) ($= 0$), the other algorithms perform poorly on BP recovery. The greedy algorithms select rows from the identity matrix that result in columns with all-zero entries and thus fail to recover most of the entries in x. Digging deeper, Insense selects rows with smaller column coherence than Convex SS and MSE-G. As a result, Insense achieves the best BP recovery performance (Table 1) among these three algorithms.

Dataset 2) Uniform/Gaussian matrix: To study the quality of the box constraint relaxation in (2), we compare Insense against the baseline algorithms for a matrix Φ where we know the globally optimal index set of rows Ω. We concatenate a 10×200 matrix with i.i.d. Gaussian entries and a 190×200 matrix with i.i.d. $[0, 1]$ uniform distribution entries. In this case, one would expect that the Gaussian submatrix has the lowest μ_{avg} when we set $M = 10$. In all 10 random trials, Insense successfully selects all Gaussian rows and hence find the globally optimal set of sensors. FrameSense and EigenMaps miss, on average, 10–20% of the Gaussian sensors. The other baselines algorithms, including Convex SS, select only a small portion (<20%) of the Gaussian rows. Table 2 indicates that Insense achieves better BP recovery performance, since it selects exclusively Gaussian rows, resulting in the minimum average coherence μ_{avg} of the resulting sensing matrix.

[2] Dashes correspond to instances where the selected matrices Φ_Ω contain columns with all zero entries; here the average coherence μ_{avg} is undefined.

Table 3: Comparison of Insense against the baseline algorithms on selecting M DNA probes to identify pathogenic samples containing K bacterial organisms.

	BP accuracy in detecting organisms %								
Number of organisms	$K = 2$			$K = 3$			$K = 5$		
Number of probes	8	12	15	12	15	20	15	20	25
Insense	**68.33**	**94.78**	**99.65**	**71.74**	**93.95**	**99.53**	**51.95**	**92.71**	**99.10**
EigenMaps	49.65	84.69	94.66	54.68	78.09	96.25	27.47	72.13	95.30
MSE-G	60.79	91.53	97.91	67.16	89.15	98.40	43.26	83.52	97.40
FrameSense	61.83	88.40	95.71	62.32	82.29	98.36	35.16	81.92	96.50
MI-G	59.98	89.68	96.40	65.69	84.10	97.39	37.96	79.72	96.00
Entropy-G	61.25	91.53	98.61	66.35	88.96	99.19	42.86	89.61	97.50
Determinant-G	46.75	82.13	94.55	48.97	76.13	96.03	24.48	72.73	92.81
Greedy SS	57.54	87.70	96.87	59.65	84.64	97.34	36.16	80.22	94.11
Convex SS	53.36	87.94	98.94	57.58	87.59	98.89	38.46	83.52	98.40
Random	61.53	88.79	96.66	62.29	86.15	97.72	38.88	82.94	86.44

Dataset 3) Microbial Diagnostics: We assess the performance of Insense on a real-world dataset from microbial diagnostics. Microbial diagnostics seek to detect and identify microbial organisms in a sample. Next-generation systems detect and classify organisms using DNA probes that bind (hybridize) to the target sequence and emit some kind of signal (e.g., fluorescence). Designing DNA probes for microbial diagnostics is an important application of sensor selection in the underdetermined sensing regime. We run Insense and the baseline sensor selection algorithms on a large sensing matrix comprising the hybridization affinity of $D = 100$ random DNA probes to $N = 42$ bacterial species (see [31]). For each algorithm, after selecting M probes and constructing a sensing matrix Φ_Ω with $|\Omega| = M$, we perform BP recovery for multiple sparse vectors x with random support. We repeat the same experiment for all $\binom{N}{K}$ sparse vectors x with $K = \{2, 3, 5\}$ non-zero elements (i.e., bacteria present) and report the average BP recovery performance on identifying the composition of the samples in Table 3. The DNA probes selected by Insense outperform all of the baseline algorithms in identifying the bacterial organisms present. Specifically, Insense requires a smaller number of DNA probes than the other algorithm to achieve almost perfect detection performance (BP accuracy $> 99\%$), suggesting that Insense is the most cost-efficient algorithm to select DNA probes for this application. Moreover, the performance gap between Insense and the other algorithms grows as the number of bacterial species present in the sample K increases, indicating that Insense has better recovery performance in complex biological samples.

5. CONCLUSIONS

In this paper, we developed the Incoherent Sensor Selection (Insense) algorithm for the underdetermined sensor selection problem that optimizes the average squared coherence of the columns of the selected sensors (rows) via a computationally efficient relaxation. Our synthetic and real-world data results have both verified the utility of the average squared coherence metric and the performance of the Insense algorithm. In particular, Insense provides superior performance than existing state-of-the-art sensor selection algorithms, especially in the real-world problem of microbial diagnostics.

6. REFERENCES

[1] S. Joshi and S. Boyd, "Sensor selection via convex optimization," *IEEE Trans. Signal Process.*, vol. 57, no. 2, pp. 451–462, Feb. 2009.

[2] M. Shamaiah, S. Banerjee, and H. Vikalo, "Greedy sensor selection: Leveraging submodularity," in *Proc. IEEE. Conf. Dec. Ctrl.*, Feb. 2010, pp. 2572–2577.

[3] J. Ranieri, A. Chebira, and M. Vetterli, "Near-optimal sensor placement for linear inverse problems," *IEEE Trans. Signal Process.*, vol. 62, no. 5, pp. 1135–1146, Mar. 2014.

[4] E. J. Candès and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 21–30, Mar. 2008.

[5] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[6] R. G. Baraniuk, "Compressive sensing," *IEEE Signal Process. Mag.*, vol. 24, no. 4, pp. 118–124, Jul. 2007.

[7] E. J. Candès, "Compressive sampling," in *Proc. Intl. Cong. Math,*, Aug. 2006, vol. 3, pp. 1433–1452.

[8] R. Tibshirani, "Regression shrinkage and selection via the Lasso," *J. Roy. Stat. Soc. B Met.*, pp. 267–288, Jan. 1996.

[9] D. L. Donoho, A. Maleki, and A. Montanari, "Message-passing algorithms for compressed sensing," *Proc. Natl. Acad. Sci.*, vol. 106, no. 45, pp. 18914–18919, 2009.

[10] J. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Trans. Inf. Theory*, vol. 50, no. 10, pp. 2231–2242, Sep. 2004.

[11] D. L. Donoho, M. Elad, and V. N. Temlyakov, "Stable recovery of sparse overcomplete representations in the presence of noise," *IEEE Trans. Inf. Theory*, vol. 52, no. 1, pp. 6–18, 2006.

[12] C. Herzet, C. Soussen, J. Idier, and R. Gribonval, "Exact recovery conditions for sparse representations with partial support information," *IEEE Trans. Inf. Theory*, vol. 59, no. 11, pp. 7509–7524, 2013.

[13] E. J. Candes, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Commun. Pur. Appl. Math.*, vol. 59, no. 8, pp. 1207–1223, Aug. 2006.

[14] R. Gribonval and P. Vandergheynst, "On the exponential convergence of matching pursuits in quasi-incoherent dictionaries," *IEEE Trans. Info. Theory*, vol. 52, no. 1, pp. 255–261, Dec. 2006.

[15] J. Ranieri, A. Vincenzi, A. Chebira, D. Atienza, and M. Vetterli, "EigenMaps: Algorithms for optimal thermal maps extraction and sensor placement on multicore processors," in *Proc. ACM/EDAC/IEEE Des. Auto. Conf. (DAC'12)*, Jun. 2012, pp. 636–641.

[16] A. Das and D. Kempe, "Algorithms for subset selection in linear regression," in *Proc. Ann. ACM Theory Comput. (STOC'08)*, May 2008, pp. 45–54.

[17] D. Golovin, M. Faulkner, and A. Krause, "Online distributed sensor selection," in *Proc. ACM Intl. Conf. Inf. Process. Sens. Net.*, Apr. 2010, pp. 220–231.

[18] A. Das and D. Kempe, "Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection," *arXiv preprint arXiv:1102.3975*, 2011.

[19] D. M. Steinberg and W. G. Hunter, "Experimental design: Review and comment," *Technometrics*, vol. 26, no. 2, pp. 71–97, May 1984.

[20] A. Krause, A. Singh, and C. Guestrin, "Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies," *J. Mach. Learn. Res.*, vol. 9, pp. 235–284, Feb. 2008.

[21] H. Wang, K. Yao, G. Pottie, and D. Estrin, "Entropy-based sensor selection heuristic for target localization," in *Proc. Intl. ACM Conf. Inf. Process. Sens. Net.*, Mar. 2004, pp. 36–45.

[22] I. Ford, D. M. Titterington, and C. P. Kitsos, "Recent advances in nonlinear experimental design," *Technometrics*, vol. 31, no. 1, pp. 49–60x, 1989.

[23] S. P. Chepuri and G. Leus, "Sparsity-promoting sensor selection for non-linear measurement models," *IEEE Trans. Signal Process.*, vol. 63, no. 3, pp. 684–698, Feb. 2015.

[24] K. B. Petersen and M. S. Pedersen, "The matrix cookbook," Tech. Rep., Technical University of Denmark, 2008.

[25] H. Attouch, J. Bolte, and B. F. Svaiter, "Convergence of descent methods for semi-algebraic and tame problems: Proximal algorithms, forward–backward splitting, and regularized Gauss–Seidel methods," *Math. Program.*, vol. 137, no. 1-2, pp. 91–129, Aug. 2013.

[26] Y. Nesterov, "Gradient methods for minimizing composite objective function," 2007.

[27] Y. Chen and X. Ye, "Projection onto a simplex," *arXiv preprint arXiv:1101.6081*, 2011.

[28] W. Wang and M. A. Carreira-Perpinán, "Projection onto the probability simplex: An efficient algorithm with a simple proof, and an application," *arXiv preprint arXiv:1309.1541*, 2013.

[29] L. Condat, "Fast projection onto the simplex and the ℓ_1-ball," 2014.

[30] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.*, vol. 20, no. 1, pp. 33–61, Aug. 1998.

[31] A. Aghazadeh, A. Y. Lin, M. A. Sheikh, A. L. Chen, L. M. Atkins, C. L. Johnson, J. F. Petrosino, R. A. Drezek, and R. G. Baraniuk, "Universal microbial diagnostics using random DNA probes," *Sci. Adv.*, vol. 2, pp. e1600025, Sep. 2016.