Homework#3: Short course on sparse recovery and compressed sensing

1. **Answer:** Three lectures about ℓ_1 minimization for sparse recovery were given today. In the first lecture, a necessary and sufficient condition for ℓ_1 recovery is firstly derived, which is determined only by the set on which the sparse signal is supported and the signs of the elements on this set. Then a more workable sufficient condition for ℓ_1 recovery is derived based on duality and optimality. Some examples are followed to show how to use these conditions in practice. In the second lecture, it is proved that if the sensing matrix is an approximate isometry for all $3S$-sparse vectors, ℓ_1 recovery is guaranteed to find the S-sparse vector. All these derivations are based on the fact that the cone condition is a sufficient condition for ℓ_1 recovery. In the last lecture, it considers how ℓ_1 recovery performs in more realistic situations, such as the signal of interest is not sparse and the measurements are noisy. It is shown that the recovery error is bounded by how well the signal is approximated by a sparse signal and the ℓ_2 norm of the measurement noise.

2. **Answer:**

   (a) The solution to this problem is the pseudo-inverse solution

   $$\hat{x} = \Phi^\dagger y = \begin{bmatrix} 4/17 \\ 1/17 \end{bmatrix}.$$

   (b) We use the sufficient condition for ℓ_1 recovery introduced in the lecture to derive the solution. For $\hat{x}$ supported on Γ and satisfying $y = \Phi\hat{x}$, $\hat{x}$ is a solution to ℓ_1 minimization if there exists $v \in \mathbb{R}$ such that

   $$\begin{aligned} (\Phi^* v)[\gamma] &= \operatorname{sgn}(\hat{x}[\gamma]), & \gamma \in \Gamma \\ |(\Phi^* v)[\gamma]| &\le 1, & \gamma \in \Gamma^c. \end{aligned}$$

   Since $\Phi^* v = \begin{bmatrix} 4v \\ v \end{bmatrix}$, it is easy to check that the second entry of $\hat{x}$ must be zero. Therefore, the constraint $y = \Phi\hat{x}$ implies that

   $$\hat{x} = \begin{bmatrix} 1/4 \\ 0 \end{bmatrix}.$$

3. **Answer:** In the simulations, the locations of the nonzero entries of the sparse vector x are randomly chosen among all possible choices, and these nonzero entries are iid Gaussian. Then x is normalized to have unit ℓ_2 norm. The sparsity S varies from 25 to 500 with step 25 and the number of observations M varies from 50 to 1000 with step 50. For each case of random matrices, the observation vector $y = \Phi x$ is calculated accordingly, and the sparse vector is estimated from Φ and y via ℓ_1 minimization. The sparse estimation is denoted as $\hat{x}$. If the recovery error $\|\hat{x} - x\|_2$ is less than 1×10^{-3}, the recovery is regarded as a success. For each choice of S and M, the simulation is repeated 100 times to calculate the probability of successful recovery.

   When the entries of the sensing matrix are iid Gaussian, the result is shown as follows, where "white" denotes 100% successful recovery and "black" denotes 0% successful recovery. From the result, it is easy to conclude how many observations we need to reliably recover an S sparse vector.

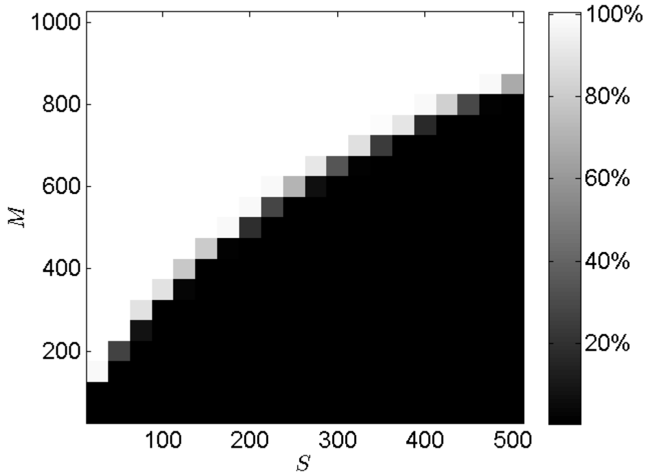

When the sensing matrix is formed by random rows of a discrete cosine transform, the result is shown as follows.

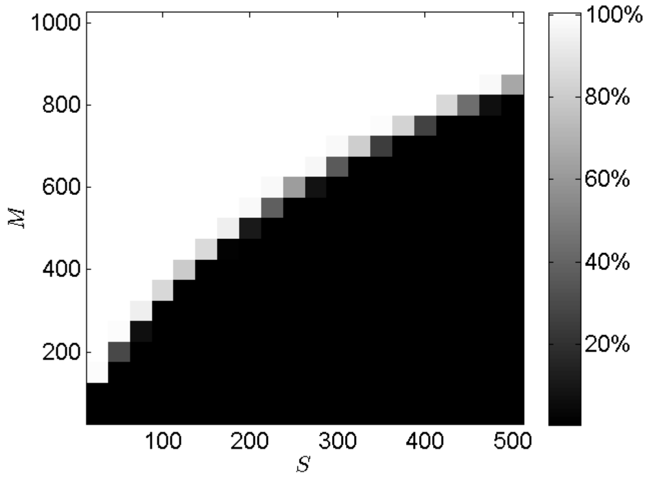

When the sensing matrix is formed by consecutive rows of a random Toeplitz matrix, the result is shown as follows.