COMS E6998-9: Algorithms for Massive Data (Spring'19) Feb 19, 2019

Lecture #9: Compressed Sensing: Restricted Isometry Property

Instructor: *Alex Andoni* Scribes: *Derrick Liu, Jie* Li

1 Review

During the last two lectures, we focused on randomized numerical linear algebra and sped up numerical linear algebra computations through sketching. We've seen two versions of linear regression: $l1$ and $l2$. There are many other problems where we can use this kind of techniques. For example in matrix multiplication where we have to compute C given some matrix A and B. such that $||C - AB||$ is small (minimizes the distance between the approximate product). This can be for different norms (2 norm, etc). Another problem is to approximate single value decomposition (SVD), where given matrix A, we want to compute U, V of rank k s.t. $||A - UV||$ is minimized. Solving this exactly and naively is trivial, however, the approaches discussed involved speeding up these computations.

2 Compressed Sensing

- **Goal** Today we will talk about compressed sensing. The problem stems from the engineering field, signal processing in particular. We are given $x \in \mathbb{R}^n$ which can be interpreted as being the 'signal'. We also think about this signal as being very large. The goal is to perform m measurements, and from the measurements we observe y. We define $y = Ax$ and hope to recover the signal x after observing y. Often times signals may have a particular structure. In our particular case, we focus on a structure x as being well approximated by a k-sparse vector. For example, you can think of a image where x is the signal in wavelet domain for Fourier domain. We know that images are sparse in a well established basis, this is why we are able to compress images. To summarize, $y = Ax$ where A is an $m \times n$ matrix, and m corresponds to the m measurements where $m \ll n$. We will show how to recover the signal with much fewer measurements through compressed sensing.

- **Motivation** To give some motivation for this type of problem, a particular interesting application for compressed sensing is the single pixel camera[**?**]. Traditionally, cameras such as a 10 mega pixel camera will have 10 million photodiodes to capture the light. Each photodiode may represent a single pixel. When you try to take an image, light will pass through the lens and photons will hit all 10 million photodiodes. The photons will then excite electrons in the photodiode, and these electrons will then be collected by the photodiodes to extract the color. The single pixel camera takes advantage of compressed sensing using a single photodiode. By taking m measurements combined with compressed sensing, it is possible to recover the image with a single photodiode, in turn saving power.

Hence, let us formalize this problem statement. Recall, we assumed that our signal x takes on a particular structure which is approximately k-sparse.

Formal problem statement:

L_0-min:	$min_{x^* \in \mathbb{R}^n} \|\|x^*\|\|_0$	(This is number of nonzeros.)
	s.t. $Ax^* = y$	$(= Ax)$

Alternative approach:

L_1-min:	$min_{x^* \in \mathbb{R}^n} \|\|x^*\|\|_1$
	s.t. $Ax^* = y$

The issue here is that the L_0-norm problem is known to be NP hard and getting the absolute minimum would require brute-force search. Instead of minimizing the L_0-norm, we will minimize the L_1-norm. This approach similarly is used in field of machine learning. We'd like to find a model that has fewest non-zeros, this is not computationally feasible, so we do the next possible thing which is minimizing the L_1 norm. Whereas minimizing L_0-norm is a non-convex problem, L_1 is the smallest convex relaxation, which can be solved as a linear program.

So we want to recover x^* (the best approximation of signal x) from y (the measurement). y = Ax. This gives us the best k-approximation. Let us define this formally. Ideally, we want:

$$x^* = \underset{x' \in \mathbb{R}^n, k-sparse}{\arg\min} ||x - x'||_1.$$

x^* is the best approximation of signal x. The solution to this arg min will be the k largest coordinates of x and zero otherwise. Similar to what we examined earlier in numerical linear algebra, we want some approximation. What we would like to see is that the relaxation x^* approximates

$$\min_{x' \in \mathbb{R}^n, x' \in k-sparse} ||x - x'||_1 = Err_1^k(x)$$

- **Some Intuition** Let us provide some intuition to understand how minimizing L_1-norm might yield solutions similar to the L_0-norm. Assume $n = 2$, we have measurement $m = 1$, lets say signal is sparse with $k = 1$. The set of possible solutions we're looking for is on the axis (see Figure 1). The dotted lines represents points with the same L_1. Let's assume we have a x that does not fall somewhere on the axis. It must lie on some line, where $y_1 = Ax'$. (See Figure 1). By finding the sparsest possible vector that lies on this line. It will be on the axis. The point that lies on this line, will still be the same x. In higher dimensions this breaks down, so it doesn't necessarily work in higher dimensions. this is purely meant to give you an intuitive sense. It is interesting to note that solving the L_1 norm here also promotes sparsity given that you are confined to the axis.

  The main theorem we will be discussing[**?**, **?**, **?**] basically states L_1-min problem solution recovers L_0-min problem solution if A is 'nice', in other words, $((k, \epsilon)-\text{RIP})$ where RIP stands for Restricted Isometry Property. Hence, understanding RIP becomes important towards compressed sensing.

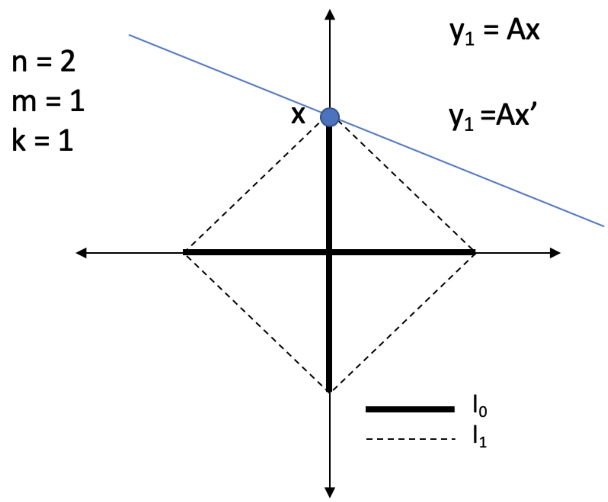

Figure 1:

Theorem 1. *Fix k; if A is i.i.d* $N(0,1)$, $m = O(k \lg \frac{n}{k})$, *then with probability* $\geq 90\%$*:*

$\forall x \in \mathbb{R}^n$, if x^* is the solution to the L_1 problem, then $||x^* - x||_1 \leq C \cdot Err_1^k(x)$.
Note 1: $C = 1 + \epsilon$, then m will depend on $1/\epsilon$
Note 2: if x is k-sparse, $x^* = x$
Note 3: x^* that is recovered from L_1 min may not be k-sparse (this is important!). However, it still is the best vector to recover. In other words, x^* does as well as any other vector to approximate.

- **Analogy to Heavy Hitters** This is basically related to count-sketch or heavy hitters in general. Remember that count-sketch was used to find heavy hitters. Heavy hitters, are finding items that occur with sufficient frequency or "heaviness", which is an approximation, for getting the most frequent items beyond a specific threshold. In the table below, we show a quick comparison to count-sketch.

Algorithm	Count-sketch	L_1 min CS
m linear measurements	O(k log n)	O(k log (n/k))
recovery time	O(n log n)	Linear Program. $n^{O(1)}$ time
probability of success	$\geq 1 - 1/n$	"deterministic"
guarantee comparison	$x^* : \|x_i^* - x_i\| \leq CErr_1^k(x)/k$	$x^* : \|\|x^* - x\|\|_1 \leq CErr_1^k(x)$

(1)

Count-sketch can be thought of as a poormans compressed sensing algorithm. In practice, the number of measurements are actually better in compressed sensing. For heavy hitters, we typically think of k as being a small value, roughly 10, for example. But here k is defined as the sparsity of the signal. This is usually the use case of compressed sensing, then log n vs log (n/k) starts to matter. Next, we introduce the Restricted Isometry Property (RIP), which plays a key role to show that solving L_1 is a good solution to L_0.

Definition 2. *$k \in N$, $\epsilon \in (0, 1/3)$, A is (k, ϵ) - Restricted Isometry Property (RIP)*

if $\forall x \in \mathbb{R}^n$ k-sparse, $||Ax||_2 = (1 \pm \epsilon)||x||_2$

To explain this definition in normal terms, what RIP desires is that we want matrix A to behave like an isometry/orthonormal matrix. If A was purely orthogonal it would be size $n \times n$, but what we want matrix A to be $m \times n$, where $m \ll n$ and still behave like it is orthonormal.

From here we will utilize two theorems to try to prove Theorem 1 we introduced earlier. The first part is to show that we can get a matrix A that satisfies the RIP with high probability. The second part is to show that once we have a matrix with this RIP property, we can solve the L_1 min problem with high probability. Hence, proving that we can recover the signal with much fewer measurements.

Theorem 3. *If A is Oblivious Subspace Embedding (O.S.E) for subspace dimension $= k$, with probability $\geq 1 - \delta$, where $\delta < \frac{0.1}{(\binom{n}{k})}$, then A is $(k, \epsilon)-RIP$*

Proof. A is O.S.E $\iff \forall U \subset \mathbb{R}^n$ is k dim: $Pr_A[\forall x \in U : ||Ax||_2 = (1 \pm \epsilon)||x||_2] \geq 1 - \delta \geq 1 - 0.1/\binom{n}{k}$
Consider U_1 ... $U_{\binom{n}{k}}$, all k-dimensional linear subspaces on k-sparse vectors. $\binom{n}{k}$ ways to choose k nonzero coordinates out of n. We can take a linear subspace, which is k dimensional, the rest are 0. Any k-sparse vectors must belong to one of these spaces. By union bound and by the fact that A is O.S.E:

$$Pr[\forall\ x \in U_1 \cup U_2 \cup ... \cup U_{\binom{n}{k}} : ||Ax||_2 = (1 \pm \epsilon)||x||_2] \geq 1 - \delta \cdot \binom{n}{k} = 0.9$$
$$\implies Pr_A[A \text{ is}(k, \epsilon) - RIP] \geq 0.9$$

□

Remark: We can construct a random A with $m = O_\epsilon(k + \lg(1/\delta)) = O(k + \lg(\binom{n}{k})) = O(k \lg(n/k))$

Theorem 4. *If A is a $(2k, \epsilon)$-RIP matrix, then $x^* = L_1$ min solution satisfies $||x^* - x||_1 \leq C \cdot Err_1^k(x)$, where $C = 1 + O(\epsilon)$*

Definition 5. *Matrix A satisfies null-space property of order k, with constant C, if $\forall \eta \in \mathbb{R}^n$ s.t. $A\eta = 0$, $\forall T \subset [n]$ of size k, we have the following:*
Denote $\eta_T = \eta$ restricted to coord T and $\eta_{-T} = \eta$ restricted to $[n] \setminus T$
$||\eta||_1 \leq C \cdot ||\eta_{-T}||_1 \implies ||\eta_T|| \leq (C - 1) \cdot ||\eta_{-T}||_1$ *since* $||\eta||_1 = ||\eta_T||_1 + ||\eta_{-T}||_1$

Lemma 6. *If matrix A satisfies $((r + 2)k, \epsilon)$-RIP, $r > 1$, then A satisfies null-space property of order $2k$ with constant $C = 1 + \sqrt{\frac{2}{r} \cdot \frac{1+\epsilon}{1-\epsilon}}$*

Proof. For η and T with $|T| = 2k$. We define sets
$T_1 = r \cdot k$ largest coordinates of η_{-T}
$T_2 = r \cdot k$ next largest coordinates
... $T_s =$ last $\leq rk$ remaining coordinates

$\hat{\eta} = \eta_T + \eta_{T_1}$

$$\begin{aligned}
&\text{Since } A\eta = 0 \\
&\qquad \implies A(\eta_T + \eta_{T_1} + ...) = 0 \\
&\qquad \implies A\hat{\eta} = A(\eta_{T_2} + \eta_{T_3} + ... + \eta_{T_s})) \\
&\qquad \implies ||\eta_T||_2 \leq ||\hat{\eta}||_2 \\
&\qquad \leq \frac{1}{1-\epsilon} \cdot ||A\hat{\eta}||_2, \text{ (applying RIP)} \\
&\qquad = \frac{1}{1-\epsilon} \cdot ||A(\eta_{T_2} + ...\eta_{T_s})||_2 \text{ (triangle inequality)} \\
&\qquad \leq \frac{1}{1-\epsilon} \sum_{j=2} ||A\eta_{T_j}||_2 \text{ (using RIP)} \\
&\qquad \leq \frac{1+\epsilon}{1-\epsilon} \sum_{j=2} ||\eta_{T_j}||_2
\end{aligned}$$

There are a few more steps that we need to finish this proof, covered in the next lecture. □