Learning from Data with Linear Algebra

Gregory Aschenbrenner

University of Connecticut

December 9th, 2020

▲□▶ ▲□▶ ▲□▶ ▲□▶ □ のQで

My Book



Is there a limit to how much we can learn from a data set? If so, how do we get there?

Outline

- 1 SVD
- 2 PCA
- 3 Norms
- 4 Numerical Linear Algebra
- 5 Randomized Linear Algebra
- 6 Changes in A^{-1} from changes in A
- 7 Interlaced Eigenvalues
- 8 Compressed Sensing and Matrix Completion

< □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □

9 Fourier Transforms

Important Ideas: The SVD

Definition:

The SVD is defined as follows $A = U\Sigma V^T$. Where *U* is an orthogonal $n \times n$ matrix, *V* is an orthogonal $m \times m$ matrix, and Σ is an $n \times m$. *U* holds our left singular vectors, *V* holds our right singular vectors and Σ holds our singular values. Note σ_1^2 to σ_r^2 are always nonzero eigenvalues of both $A^T A$ and AA^T .

< ロ > < 同 > < 三 > < 三 > < 三 > < ○ </p>



▲□▶ ▲□▶ ▲目▶ ▲目▶ 三日 - 釣A@

Quick example of the SVD

$$A = \begin{bmatrix} 6 & 2 \\ 2 & 4 \\ 7 & 8 \end{bmatrix} AA^{T} = \begin{bmatrix} 40 & 20 & 58 \\ 20 & 20 & 46 \\ 58 & 46 & 113 \end{bmatrix} A^{T}A = \begin{bmatrix} 89 & 76 \\ 76 & 84 \end{bmatrix}$$

$$\begin{split} \Sigma &= \begin{bmatrix} 12.74916 & 0 \\ 0 & 3.23402113414883 \\ 0 & 0 \end{bmatrix} \\ U &= \begin{bmatrix} -0.44729 & 0.84570 & -0.29104 \\ -0.33090 & -0.45880 & -0.82462 \\ -0.83092 & -0.27253 & 0.48507 \end{bmatrix} \\ V^T &= \begin{bmatrix} -0.71863 & -0.69538 \\ 0.69538 & -0.71863 \end{bmatrix} \end{split}$$

Important Ideas: Principal Component Analysis

Now that we have the SVD, what can it do?

The notion of PCA is to take the infomation available to us from the SVD and pick what infomation is important.

The real focus of utilizing PCA is to find the important pieces of the matrix use these to make determinations about data sets. In this way PCA can be thought of as unguided learning in some sense, with only linear algebra guiding us.

・ロト ・ 理 ト ・ ヨ ト ・ ヨ ・ うゅつ

Important Ideas: Principal Component Analysis

What is the best rank k matrix, A_k that approximates A?

Using the idea of the SVD we define this matrix $A_k = \sigma_1 u_1 v_1^T + ... + \sigma_k u_k v_k^T$. Notice the matrix is composed of k rank 1 matrices.

・ロト ・ 理 ト ・ ヨ ト ・ ヨ ・ うゅつ

Eckart-Young

Thm: If B has rank k then $||A - B|| \ge ||A - A_k||$

Example of PCA

Using the A in our SVD example,
$$A = \begin{bmatrix} 6 & 2 \\ 2 & 4 \\ 7 & 8 \end{bmatrix}$$
 What is the best
 A_k to approximate A? It is $A_k = \sigma_1 u_1 v_1^T + \dots + \sigma_k u_k v_k^T$! Thus
the best rank 2 matrix that approximates would be
 $A_2 = \sigma_1 u_1 v_1^T + \sigma_2 u_2 v_2^T = \begin{bmatrix} 12.7492 \\ 3.2340 \end{bmatrix}$

◆□▶ ◆□▶ ◆ □▶ ◆ □▶ ● □ ● ● ● ●



(a) PCA on wider spread of data

(b) PCA on tightly grouped data

・ロット (雪) (日) (日)

э.

Figure: PCA in two very different data sets

So many norms so little time!

What is a norm?

A norm of a matrix, as we see below, is really a measure of the size of the elements of the matrix.

Types of norms useful here

Spectral Norm:

$$\|\boldsymbol{A}\|_{2} = max \frac{\|\boldsymbol{A}x\|}{\|\boldsymbol{x}\|} = \sigma_{1}$$

Forbenius Norm:

$$\|\boldsymbol{A}\|_{\boldsymbol{F}} = \sqrt{\sigma_1^2 + \ldots + \sigma_r^2}$$

Nuclear Norm:

$$\|\boldsymbol{A}\|_{\boldsymbol{N}} = \sigma_1 + \dots + \sigma_r$$

What happens when our matrix is poorly behaved?

Three main cases: A is square and singular, ill-conditioned, or simply incredibly massive.

To solve these issues we use a variety of methods, including interesting concepts such as recusive and dynamic least squares, Numerical Linear Algebra, Randomized Linear Algebra, as well as manipulation of the column space's basis to yield an easier problem.

・ロ・・聞・・ヨ・・ヨ・ うくつ

Methods we are interested in

- Numerical Linear Algebra
- Randomized Linear Algebra

Numerical Linear Algebra

Orthogonal or bust!

For our purposes we will utilize iterative techniques from Numerically Linear Algebra to orthogonalize a matrix of various sizes.

Methods Used

- Arnoldi Iterations
- Lanczos Iterations
- Conjugate Gradient

Definition:

Krylov Subspace: Given A and b, we can compute $b, Ab, ...A^{n-1}b$, the combination of these n vectors produce a *n*th Krylov Subspace.

Arnoldi Iteration

<pre>1 Gunction [Q,H] = arnoldi(A,ql,m) 2 3</pre>		
7 - H = zeros(min(m+1,m), <u>n</u>); 8 9 - for k=1:m 10 - $z = \lambda^{*}Q(:, k);$	Arnoldi Iteration $v = Aq_k$	$m{q}_1 = m{b}/ m{b} , m{q}_2, \dots, m{q}_k$ are known Start with new $m{v}$
$\begin{array}{cccc} 11 - & & \text{for } i=1:k \\ 12 - & & H(i,k) = Q(:,i) \\ 3 - & & z = F(i,k) \partial Q(:,i) \\ \end{array}$	for $j = 1$ to k	For each known q
$\begin{array}{rcl} 14 - & \text{end} \\ 15 - & \text{if } k < n \end{array}$	$h_{jk} = \boldsymbol{q}_j^{\scriptscriptstyle \perp} \boldsymbol{v}$ $\boldsymbol{v} = \boldsymbol{v} - h_{jk} \boldsymbol{q}_j$	Subtract projection
$ \begin{array}{ll} 16 & - & n_1(k^{+1}, k) = n_2(m_1(2)); \\ 17 & - & \text{if } H_1(k^{+1}, k) = 0, \text{ return, end} \\ 18 & - & Q(;, k^{+1}) = z/H(k^{+1}, k); \\ 19 & - & \text{end} \end{array} $	$egin{aligned} h_{k+1,k} &= oldsymbol{v} \ oldsymbol{q}_{k+1} &= oldsymbol{v}/h_{k+1,k} \end{aligned}$	Compute norm New basis vector with norm 1

(a) MATLAB Code

(b) Pseudocode

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ - 三 - のへぐ

Figure: Implementation of Arnoldi Iterations

Given:

$$A = \begin{bmatrix} 3 & 2 & 4 \\ 1 & 6 & 3 \\ 0 & 0 & 300 \end{bmatrix}$$

Let
$$q_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = \frac{b}{\|b\|}$$
 we recieve what we expect, $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$

Let
$$q_1 = \begin{bmatrix} 3 \\ 6 \\ \pi \end{bmatrix} = \frac{b}{\|b\|}$$
 yields, $\begin{bmatrix} 0.4050 & -0.1806 & -0.8963 \\ 0.8100 & -0.3838 & 0.4434 \\ 0.4241 & 0.9056 & 0.0091 \end{bmatrix}$

Arnoldi Iterations on vector b and matix A yields $Q_k^T A Q_k = H_k$ where H_k is Hessenberg Matrix

Hessenberg Matrix is special square matrix, that is close to triangular. Upper Hessenberg has zero entries below subdiagonal; lower Hessenverg has zero entries above the first superdiagonal.

	[1	4	2	3]
1 _	3	4	1	7
A —	0	2	3	4
	0	0	1	3
	1	2	0	0]
P	5	2	3	0
D =	3	4	3	7
	5	6	1	1

Lanczos Iteration

Really is just Arnoldi Iteration for a symmetric matrice. This simplifies the process.

28	% Initialize variables
29	Q = nan(n, k);
30	q = v / norm(v);
31	Q(:, 1) = q;
32	d = nan(k, 1);
33	od = nan(k-1,1);
34	% Perform Lanczos iterations
35	for i = 1:k
36	z = mat * q;
37	d(i) = q* * z;
38	z = z - Q(:,1:i) * (Q(:,1:i)' * z);
39	
40	if (i ~= k)
41	od(i) = norm(z);
42	q = z / od(i);
43	Q(:, i + 1) = q;
44	end
45	- end
46	& Construct T
47	T = diag(d) + diag(od, -1) + diag(od, 1);
48	% Return user-requested information
49	if (nargout 2)
50	<pre>varargout(1) = Q;</pre>
51	- end
52	

Lanczos iteration for $Sx = \lambda x$ (symmetric Arnoldi)				
$q_0 = 0, \ q_1 = b/ b $	Orthogonalize b. Sb, Sb ²			
For $k = 1, 2, 3, \ldots$				
$v = Sq_k$	Start with new v			
$a_k = \boldsymbol{q}_k^{\mathrm{T}} \boldsymbol{v}$	Diagonal entry in T is a_k			
$\boldsymbol{v} = \boldsymbol{v} - b_{k-1}\boldsymbol{q}_{k-1} - a_k \boldsymbol{q}_k$	Orthogonal to earlier q's			
$b_k = \boldsymbol{v} $	Off-diagonal entry in T is b_k			
$\boldsymbol{q}_{k+1} = \boldsymbol{v}/b_k$	Next basis vector			

(a) MATLAB Code

(b) Pseudocode

Figure: Implementation of Lanczos Iterations

Conjugate Gradient Iteration

Solves Sx = b when S is positive definite.

1	<pre>function x = conjgrad(A,b,tol)</pre>
2	if nargin<3
3	tol=1e-10;
4	end
5	x = b;
6	$r = b - A^*x;$
7	if norm(r) < tol
8	return
9	end
10	y = -r;
11	z = A*y;
12	s = y'*z;
13	t = (r'*y)/s;
14	$x = x + t^*y;$
15	
16	for k = 1:numel(b);
17	$r = r - t^{*}z;$
18	<pre>if(norm(r) < tol)</pre>
19	return;
20	end
21	B = (r'*z)/s;
22	y = -r + B*y;
23	$z = A^*y;$
24	s = y'*z;
25	t = (r'*y)/s;
26	$x = x + t^*y;$
27	- end
28	end

Conjugate Gradient Iteration for Positive Definite S $x_0 = 0, r_0 = b, d_0 = r_0$

for k = 1 to N $\alpha_k = (\boldsymbol{r}_{k-1}^{\mathrm{T}} \boldsymbol{r}_{k-1})/(\boldsymbol{d}_{k-1}^{\mathrm{T}} \boldsymbol{S} \boldsymbol{d}_{k-1})$ step length \boldsymbol{x}_{k-1} to \boldsymbol{x}_k $\boldsymbol{x}_k = \boldsymbol{x}_{k-1} + \alpha_k \boldsymbol{d}_{k-1}$ $\boldsymbol{r}_k = \boldsymbol{r}_{k-1} - \alpha_k S \boldsymbol{d}_{k-1}$ $\beta_k = (\boldsymbol{r}_k^{\mathrm{T}} \boldsymbol{r}_k) / (\boldsymbol{r}_{k-1}^{\mathrm{T}} \boldsymbol{r}_{k-1})$ $d_k = r_k + \beta_k d_{k-1}$

approximate solution new residual $b - Sx_k$ improvement this step next search direction

% Notice : only 1 matrix-vector multiplication Sd in each step

(a) MATLAB Code

(b) Pseudocode

Figure: Implementation of Conjugate Gradient

Randomized Linear Algebra

Idea:

Make random vector x have one element x_k s.t. Ax is a random sample of the column space of A. Let S be a random sample matrix with columns composed of x vectors.

・ロト ・四ト ・ヨト ・ヨト

Let A and B be arbitary matrices

C = AS and $R = S^T B$ and $CR = ASS^T B \approx AB$

Randomized Linear Algebra

How to keep optimize further?

We can use probability! Specifically mean and variance.

S is n by s sampling matrix, with one nonzero value per column. Let $A = [a_1 a_2 a_3]$.

• 1 See that
$$AS = [a_1, a_2, a_3] \begin{bmatrix} s_{11} & 0 \\ 0 & 0 \\ 0 & s_{32} \end{bmatrix} = \begin{bmatrix} s_{11}a_1 & s_{32}a_3 \end{bmatrix}$$
.
Now must choose values of s_{ki} , but how?

• 2 Assign probability p_j to all n columns of A, with $p_1 + p_2 + ... + p_n = 1$

 3 Choose S columns with replacement (can choose more than once)

Randomized Linear Algebra

- 4 If column k of A (row k of B) is choosen multipy both by $\frac{1}{\sqrt{SD_k}}$
- 5 (col k of A)(row k of B)/spk goes into random product AB

So far just the mean is correct, can we use norms to get a correct variance?

Yes, but they are technical to derive! (So we won't)

Mean: $E[X] = \frac{1}{s}AB$ Variance: $E[||AB - CR||_F^2] = \frac{1}{s}(C^2 - ||AB||_F^2)$

Changes in A^{-1} from changes in A

Easing into it

$$M = I - uv^{T} \rightarrow M^{-1} = I + \frac{uv^{T}}{1 - v^{T}u}$$

Correction term is rank 1. If $v^T u = 1$ then M^{-1} does not exist.

V is $n \times k$, U is also $n \times k$

$$M = I_n - UV^T \to M^{-1} = I_n + U(I_k - V^T U)^{-1} V^T$$

Sherman-Morrison-Woodbury Formula

$$M^{-1} = (A - UV^{T})^{-1} = A^{-1} + A^{-1}U(I - V^{T}A^{-1}U)^{-1}V^{T}A^{-1}$$

◆□▶ ◆□▶ ◆三▶ ◆三▶ ・三 ・ つへの

Derivatives of A^{-1}

Let $B = A + \Delta A$

$$B^{-1} - A^{-1} = B^{-1}(A - B)A^{-1} \rightarrow \frac{\Delta A^{-1}}{\Delta t} = -(A + \Delta A)^{-1}\frac{\Delta A}{\Delta t}A^{-1}$$

Approaches $\frac{dA^{-1}}{dt} = -A^{-1}\frac{dA}{dt}A^{-1}$

Applications

- The Kalman Filter, is the updating of dynamic least squares. Thus even when there is no new data the state vector x changes with time. Think GPS and the time between new data and what is happening on the ground.
- Quasi-Newton Update Methods, updates the Jacobian Matrix in the classical Newton Method instead of recalculating.

Interlacing Eigenvalues

Fundamental Question: How does each λ change as A changes?

Interesting anwser, as it is matters what time period we are observing change over. When we take the derivative, $\frac{d\lambda}{dt}$ we have minimal problems are the derivative is a linear operator. But what about $\lambda(A + \Delta A)$? Turns out this is a very difficult question to answer.

Let S be changed to $S + uu^T$

We say uu^T is positive semidefinite $(x^T M x \ge 0 \in \mathbb{R}^n)$. This addition only increases the value of the eigenvalues, thus the change in the eigenvalues is $\lambda_1 \ge \lambda_2 \ge \dots$ to $z_1 \ge z_2 \ge \dots$.

Interlacing Eigenvalues

Each z_i of $S + uu^T$ is not smaller than λ_i or greater than λ_{i-1} . This idea forces the relation $z_1 \ge \lambda_1 \ge z_2 \ge \lambda_2 \ge ... \ge z_n \ge \lambda_n$.

・ロト・日本・ビア・ビア・ ビー うぐつ

Derivative of an Eigenvalue

Given a time dependent A

We start with formula $A(t)x(t) = \lambda(t)x(t)$ simply multiply by y^T and use $y^T x = 1$. Yields $\lambda = y^T A x$, this the derivative and cancelling terms yields $\frac{d\lambda}{dt} = y^T \frac{dA}{dt} x$.

When the change to S is $\theta u u^T$ what happens?

Actually yields the Secular Equation, $\frac{1}{\theta} = u^T (zI - S)^{-1} u = \sum_{k=1}^n \frac{c_k^2}{z - \lambda_k}$. z are the n eigenvalues of $S + \theta u u^T$.

◆□▶ ◆□▶ ◆臣▶ ◆臣▶ 臣 の�?

z_k doesn't go past λ_{k-1}



◆□▶ ◆□▶ ◆三▶ ◆三▶ ・三 のへで

Compressed Sensing and Matrix Completion

Idea behind Compressed Sensing is taking a sparse signal and completeing it to its full state from incompleted data

For Compressed Sensing to function properly the basis, V, must use as few v_n as possible. Further the signal must have another basis, W to represent it. For Compressed Sensing there must be incoherence of V and W, or entries of $V^T W$ are small. Common choice for V and W is F (Fourier Matrix) and I, as the entries of F have equal size.

Matrix Completion seeks to take an incomplete matrix A_0 and complete it to A while keeping rank low as possible

More formally seeks to minimize $||A||_N$ subject to $A = A_0$ in the known entries.

Matrix Completion

Say we know K entries in $n \times n$ matrix of rank r, can we perfectly recover the rest of the data?

While it may seem odd, yes we can. Perfect recovery of A is highly probable! Must be careful as we can force failures.

・ロマ・山 マ・山 マ・山 マ・山 マショー

Fourier Transforms

What is a Fourier Transform?

Real series	$f(x) = a_0 + a_1 \cos x + b_1 \sin x + a_2 \cos 2x + b_2 \sin 2x + \cdots$
Complex series	$f(x) = c_0 + c_1 e^{ix} + c_{-1} e^{-ix} + c_2 e^{2ix} + c_{-2} e^{-2ix} + \cdots$
Fourier integrals	$oldsymbol{f}(oldsymbol{x}) = \int_{-\infty}^{\infty} \widehat{f}(k) e^{ikx} dk$
Discrete series	$f = c_0 b_0 + c_1 b_1 + \dots + c_{N-1} b_{N-1}$ = Fourier matrix F times c

The reason all of the coefficents have nice formula is due to orthogonality!

< □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □

Fourier Transforms

Fourier Transform Matrix F and Discrete Fourier Transform Matrix $\boldsymbol{\Omega}$

$$F_{4} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & i & i^{2} & i^{3} \\ 1 & j^{2} & i^{4} & j^{6} \\ 1 & i^{3} & i^{6} & j^{9} \end{bmatrix}$$
$$\Omega_{4} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -i & (-i)^{2} & (-i)^{3} \\ 1 & (-i)^{2} & (-i)^{4} & (-i)^{6} \\ 1 & (-i)^{3} & (-i)^{6} & (-i)^{9} \end{bmatrix}$$
$$F\Omega = NI$$



Strang, G. (2019). Linear algebra and learning from data. Wellesley, MA: Wellesley-Cambridge Press.

