

LINEAR ALGEBRA AND LEARNING FROM DATA

GILBERT STRANG

Massachusetts Institute of Technology

WELLESLEY PUBLISHERS - INDIA

Linear Algebra and Learning from Data

Copyright ©2020 by Gilbert Strang

ISBN 978-1-7331466-2-3

All rights reserved. No part of this book may be reproduced or stored or transmitted by any means, including photocopying, without written permission from Wellesley - Cambridge Press. Translation in any language is strictly prohibited. Authorized translations are arranged by the publisher.

L^AT_EX typesetting by **Ashley C. Fernandes** (info@problemsolvingpathway.com)

Printed in India

9 8 7 6 5 4 3 2 1

Other texts from Wellesley - Cambridge Press

Introduction to Linear Algebra, 5th Edition , Gilbert Strang	ISBN 978-0-9802327-7-6
Computational Science and Engineering , Gilbert Strang	ISBN 978-0-9614088-1-7
Wavelets and Filter Banks , Gilbert Strang and Truong Nguyen	ISBN 978-0-9614088-7-9
Introduction to Applied Mathematics , Gilbert Strang	ISBN 978-0-9614088-0-0
Calculus, Third edition (2017) , Gilbert Strang	ISBN 978-0-9802327-5-2
Algorithms for Global Positioning , Kai Borre & Gilbert Strang	ISBN 978-0-9802327-3-8
Essays in Linear Algebra , Gilbert Strang	ISBN 978-0-9802327-6-9
Differential Equations and Linear Algebra , Gilbert Strang	ISBN 978-0-9802327-9-0
An Analysis of the Finite Element Method , 2017 edition, Gilbert Strang and George Fix	ISBN 978-0-9802327-8-3

Wellesley Publishers, India
Gilbert Strang@MIT

www.wellesleypublishers.com/buy.html
math.mit.edu/~gs

The website for this book is math.mit.edu/learningfromdata

That site will link to 18.065 course material and video lectures on YouTube and OCW

The cover photograph shows a neural net on Inle Lake. It was taken in Myanmar.

From that photograph Lois Sellers designed and created the cover.

The snapshot of playground.tensorflow.org was a gift from its creator Daniel Smilkov.

Linear Algebra is included in MIT's OpenCourseWare site ocw.mit.edu

This provides video lectures of the full linear algebra course 18.06 and 18.06 SC

This book can be sold and purchased only in India

Preface and Acknowledgments

My deepest gratitude goes to Professor Raj Rao Nadakuditi of the University of Michigan. On his sabbatical in 2017, Raj brought his EECS 551 course to MIT. He flew to Boston every week to teach 18.065. Thanks to Raj, the students could take a new course. He led the in-class computing, he assigned homeworks, and exams were outlawed.

This was linear algebra for signals and data, and it was alive. 140 MIT students signed up. Alan Edelman introduced the powerful language *Julia*, and I explained the four fundamental subspaces and the Singular Value Decomposition. The labs from Michigan involved rank and SVD and applications. We were asking the class for *computational thinking*.

That course worked, even the first time. It didn't touch one big topic : **Deep learning**. By this I mean the excitement of creating a learning function on a neural net, with the hidden layers and the nonlinear activation functions that make it so powerful. The system trains itself on data which has been correctly classified in advance. The optimization of weights discovers important features—the shape of a letter, the edges in an image, the syntax of a sentence, the identifying details of a signal. Those features get heavier weights—*without overfitting the data* and learning everything. Then unseen test data from a similar population can be identified by virtue of having those same features.

The algorithms to do all that are continually improving. Better if I say that they are *being improved*. This is the contribution of computer scientists and engineers and biologists and linguists and mathematicians and especially optimists—those who can optimize weights to minimize errors, and also those who believe that deep learning can help in our lives.

You can see why this book was written :

1. To organize central methods and ideas of **data science**.
2. To see how the language of **linear algebra** gives expression to those ideas.
3. Above all, to show how to **explain and teach** those ideas—to yourself or to a class.

I certainly learned that projects are far better than exams. Students ask their own questions and write their own programs. From now on, **projects!**

Linear Algebra and Calculus

The reader will have met the two central subjects of undergraduate mathematics: Linear algebra and calculus. For deep learning, it is linear algebra that matters most. We compute “weights” that pick out the important features of the training data, and those weights go into matrices. The form of the learning function is described on page iv. Then calculus shows us the *direction to move*, in order to improve the current weights x_k .

From calculus, it is partial derivatives that we need (and not integrals):

Reduce the error $L(x)$ by moving from x_k to $x_{k+1} = x_k - s_k \nabla L$.

That symbol ∇L stands for the first derivatives of $L(x)$. Because of the minus sign, x_{k+1} is *downhill* from x_k on the graph of $L(x)$. The stepsize s_k (also called the learning rate) decides how far to move. You see the basic idea: Reduce the loss function $L(x)$ by moving in the direction of fastest decrease. $\nabla L = \mathbf{0}$ at the best weights x^* .

The complication is that the vector x represents thousands of weights. So we have to compute thousands of partial derivatives of L . And L itself is a complicated function depending on several layers of x 's as well as the data. So we need the chain rule to find ∇L .

The introduction to Chapter VI will recall essential facts of multivariable calculus.

By contrast, *linear algebra is everywhere in the world of learning from data*. This is the subject to know! The first chapters of this book are essentially a course on applied linear algebra—the basic theory and its use in computations. I can try to outline how that approach (to the ideas we need) compares to earlier linear algebra courses. Those are quite different, which means that there are good things to learn.

Basic course

1. Elimination to solve $Ax = b$
2. Matrix operations and inverses and determinants
3. Vector spaces and subspaces
4. Independence, dimension, rank of a matrix
5. Eigenvalues and eigenvectors

If a course is mostly learning definitions, that is not linear algebra in action. A stronger course puts the algebra to use. The definitions have a purpose, and so does the book.

Stronger course

1. $Ax = b$ in all cases: square system—too many equations—too many unknowns.
2. Factor A into LU and QR and $U\Sigma V^T$ and CMR : *Columns times rows*.
3. *Four fundamental subspaces*: dimensions and orthogonality and good bases.
4. Diagonalizing A by eigenvectors and by left and right singular vectors.
5. Applications: Graphs, convolutions, iterations, covariances, projections, filters, networks, images, matrices of data.

Linear algebra has moved to the center of machine learning, and we need to be there.

A book was needed for the 18.065 course. It was started in the original 2017 class, and a first version went out to the 2018 class. I happily acknowledge that this book owes its existence to Ashley C. Fernandes. Ashley receives pages scanned from Boston and sends back new sections from Mumbai, ready for more work. This is our seventh book together and I am extremely grateful.

Students were generous in helping with both classes, especially William Loucks and Claire Khodadad and Alex LeNail and Jack Strang. The project from Alex led to his online code alexlenail.me/NN-SVG/ to draw neural nets (an example appears on page v). The project from Jack on <http://www.teachyourmachine.com> learns to recognize hand-written numbers and letters drawn by the user: open for experiment. And the new site math.mit.edu/ENNUI aims to help simplify the construction of a deep learning function.

MIT's faculty and staff have given generous and much needed help:

Suvrit Sra gave a fantastic lecture on stochastic gradient descent (now an 18.065 video)

Alex Postnikov explained when matrix completion can lead to rank one (Section IV.8)

Tommy Poggio showed his class how deep learning generalizes to new data

Jonathan Harmon and Tom Mullaly and Liang Wang contributed to this book every day

Ideas arrived from all directions and gradually they filled this textbook.

The Content of the Book

This book aims to explain the mathematics on which data science depends: *Linear algebra, optimization, probability and statistics*. The weights in the learning function go into matrices. Those weights are optimized by “stochastic gradient descent”. That word stochastic (= random) is a signal that success is governed by probability not certainty. The law of large numbers extends to the law of large functions: If the architecture is well designed and the parameters are well computed, there is a high probability of success.

Please note that this is not a book about computing, or coding, or software. Many books do those parts well. One of our favorites is *Hands-On Machine Learning* (2017) by Aurélien Géron (published by O'Reilly). And online help, from Tensorflow and Keras and MathWorks and Caffe and many more, is an important contribution to data science.

Linear algebra has a wonderful variety of matrices: symmetric, orthogonal, triangular, banded, permutations and projections and circulants. In my experience, *positive definite symmetric matrices* S are the aces. They have positive eigenvalues λ and orthogonal eigenvectors \mathbf{q} . They are combinations $S = \lambda_1 \mathbf{q}_1 \mathbf{q}_1^T + \lambda_2 \mathbf{q}_2 \mathbf{q}_2^T + \dots$ of simple rank-one projections $\mathbf{q} \mathbf{q}^T$ onto those eigenvectors. And if $\lambda_1 \geq \lambda_2 \geq \dots$ then $\lambda_1 \mathbf{q}_1 \mathbf{q}_1^T$ is the most informative part of S . For a sample covariance matrix, that part has the greatest variance.

Chapter I In our lifetimes, the most important step has been to extend those ideas from symmetric matrices to all matrices. Now we need **two sets of singular vectors, u 's and v 's**. Singular values σ replace eigenvalues λ . The decomposition $A = \sigma_1 u_1 v_1^T + \sigma_2 u_2 v_2^T + \dots$ remains correct (this is the SVD). With decreasing σ 's, those rank-one pieces of A still come in order of importance. That “Eckart-Young Theorem” about A complements what we have long known about the symmetric matrix $A^T A$: For rank k , stop at $\sigma_k u_k v_k^T$.

II The ideas in Chapter I become algorithms in Chapter II. For quite large matrices, the σ 's and u 's and v 's are computable. For very large matrices, we resort to randomization: Sample the columns and the rows. For wide classes of big matrices this works well.

III-IV Chapter III focuses on low rank matrices, and Chapter IV on many important examples. We are looking for properties that make the computations especially fast (in III) or especially useful (in IV). The *Fourier matrix* is fundamental for every problem with constant coefficients (not changing with position). That discrete transform is superfast because of the **FFT**: the Fast Fourier Transform.

V Chapter V explains, as simply as possible, the statistics we need. The central ideas are always **mean and variance**: The *average* and the *spread* around that average. Usually we can reduce the mean to zero by a simple shift. Reducing the variance (the uncertainty) is the real problem. For random vectors and matrices and tensors, that problem becomes deeper. It is understood that the *linear algebra of statistics* is essential to machine learning.

VI Chapter VI presents two types of optimization problems. First come the nice problems of linear and quadratic programming and game theory. Duality and saddle points are key ideas. But the goals of deep learning and of this book are elsewhere: *Very large problems with a structure that is as simple as possible*. “Derivative equals zero” is still the fundamental equation. The second derivatives that Newton would have used are too numerous and too complicated to compute. Even using all the data (when we take a descent step to reduce the loss) is often impossible. That is why we choose only a minibatch of input data, in each step of stochastic gradient descent.

The success of large scale learning comes from the wonderful fact that *randomization often produces reliability*—when there are thousands or millions of variables.

VII Chapter VII begins with the architecture of a neural net. An input layer is connected to hidden layers and finally to the output layer. For the training data, input vectors v are known. Also the correct outputs are known (often w is the correct classification of v). **We optimize the weights x in the learning function F so that $F(x, v)$ is close to w for almost every training input v .**

Then F is applied to *test data*, drawn from the same population as the training data. If F learned what it needs (*without overfitting*: we don't want to fit 100 points by 99th degree polynomials), the test error will also be low. The system recognizes images and speech. It translates between languages. It may follow designs like ImageNet or AlexNet, winners of major competitions. A neural net defeated the world champion at Go.

The function F is often *piecewise linear*—the weights go into matrix multiplications. Every neuron on every hidden layer also has a nonlinear “activation function”. The ramp function $\text{ReLU}(x) = (\text{maximum of } 0 \text{ and } x)$ is now the overwhelming favorite.

There is a growing world of expertise in designing the layers that make up $F(\mathbf{x}, \mathbf{v})$. We start with *fully connected* layers—all neurons on layer n connected to all neurons on layer $n + 1$. Often CNN’s are better—*Convolutional neural nets* repeat the same weights around all pixels in an image: a very important construction. Other layers are different. A *pooling layer* reduces the dimension. *Dropout* randomly leaves out neurons. *Batch normalization* resets the mean and variance. All these steps create a function that closely matches the training data. Then $F(\mathbf{x}, \mathbf{v})$ is ready to use.

Acknowledgments

Above all, I welcome this chance to thank so many generous and encouraging friends :

Pawan Kumar and Leonard Berrada and Mike Giles and Nick Trefethen in Oxford
 Ding-Xuan Zhou and Yunwen Lei in Hong Kong
 Alex Townsend and Heather Wilber at Cornell
 Nati Srebro and Srinadh Bhojanapalli in Chicago
 Tammy Kolda and Thomas Strohmer and Trevor Hastie and Jay Kuo in California
 Bill Hager and Mark Embree and Wotao Yin, for help with Chapter III
 Stephen Boyd and Lieven Vandenbergh, for great books
 Alex Strang, for creating the best figures, and more
 Ben Recht in Berkeley, especially.

Your papers and emails and lectures and advice were wonderful.

THE MATRIX ALPHABET

A	Any Matrix	Q	Orthogonal Matrix
C	Circulant Matrix	R	Upper Triangular Matrix
C	Matrix of Columns	R	Matrix of Rows
D	Diagonal Matrix	S	Symmetric Matrix
F	Fourier Matrix	S	Sample Covariance Matrix
I	Identity Matrix	T	Tensor
L	Lower Triangular Matrix	U	Upper Triangular Matrix
L	Laplacian Matrix	U	Left Singular Vectors
M	Mixing Matrix	V	Right Singular Vectors
M	Markov Matrix	X	Eigenvector Matrix
P	Probability Matrix	Λ	Eigenvalue Matrix
P	Projection Matrix	Σ	Singular Value Matrix

Video lectures : OpenCourseWare ocw.mit.edu and YouTube (**Math 18.06 and 18.065**)

Introduction to Linear Algebra (5th ed) by Gilbert Strang, Wellesley-Cambridge Press

Book websites : math.mit.edu/linearalgebra and math.mit.edu/learningfromdata

Table of Contents

Deep Learning and Neural Nets	iii
Preface and Acknowledgments	vi
Part I : Highlights of Linear Algebra	1
I.1 Multiplication Ax Using Columns of A	2
I.2 Matrix-Matrix Multiplication AB	9
I.3 The Four Fundamental Subspaces	14
I.4 Elimination and $A = LU$	21
I.5 Orthogonal Matrices and Subspaces	29
I.6 Eigenvalues and Eigenvectors	36
I.7 Symmetric Positive Definite Matrices	44
I.8 Singular Values and Singular Vectors in the SVD	56
I.9 Principal Components and the Best Low Rank Matrix	71
I.10 Rayleigh Quotients and Generalized Eigenvalues	81
I.11 Norms of Vectors and Functions and Matrices	88
I.12 Factoring Matrices and Tensors : Positive and Sparse	97
Part II : Computations with Large Matrices	113
II.1 Numerical Linear Algebra	115
II.2 Least Squares : Four Ways	124
II.3 Three Bases for the Column Space	138
II.4 Randomized Linear Algebra	146

Part III: Low Rank and Compressed Sensing	159
III.1 Changes in A^{-1} from Changes in A	160
III.2 Interlacing Eigenvalues and Low Rank Signals	168
III.3 Rapidly Decaying Singular Values	178
III.4 Split Algorithms for $\ell^2 + \ell^1$	184
III.5 Compressed Sensing and Matrix Completion	195
Part IV: Special Matrices	203
IV.1 Fourier Transforms : Discrete and Continuous	204
IV.2 Shift Matrices and Circulant Matrices	213
IV.3 The Kronecker Product $A \otimes B$	221
IV.4 Sine and Cosine Transforms from Kronecker Sums	228
IV.5 Toeplitz Matrices and Shift Invariant Filters	232
IV.6 Graphs and Laplacians and Kirchoff's Laws	239
IV.7 Clustering by Spectral Methods and k -means	245
IV.8 Completing Rank One Matrices	255
IV.9 The Orthogonal Procrustes Problem	257
IV.10 Distance Matrices	259
Part V: Probability and Statistics	263
V.1 Mean, Variance, and Probability	264
V.2 Probability Distributions	275
V.3 Moments, Cumulants, and Inequalities of Statistics	284
V.4 Covariance Matrices and Joint Probabilities	294
V.5 Multivariate Gaussian and Weighted Least Squares	304
V.6 Markov Chains	311

Part VI: Optimization	321
VI.1 Minimum Problems : Convexity and Newton’s Method	324
VI.2 Lagrange Multipliers = Derivatives of the Cost	333
VI.3 Linear Programming, Game Theory, and Duality	338
VI.4 Gradient Descent Toward the Minimum	344
VI.5 Stochastic Gradient Descent and ADAM	359
Part VII: Learning from Data	371
VII.1 The Construction of Deep Neural Networks	375
VII.2 Convolutional Neural Nets	387
VII.3 Backpropagation and the Chain Rule	397
VII.4 Hyperparameters : The Fateful Decisions	407
VII.5 The World of Machine Learning	413
Books on Machine Learning	416
Eigenvalues and Singular Values : Rank One	417
Codes and Algorithms for Numerical Linear Algebra	418
Counting Parameters in the Basic Factorizations	419
Index of Authors	420
Index	423
Index of Symbols	432