

ST758: Computation for Statistical Research

Tue/Thu 10:15am-11:30am, SAS Hall 1108

Instructor: Dr Hua Zhou, hua_zhou@ncsu.edu


1 Lecture 1: Aug 21

Today

- Introduction and course logistics
- Computer storage and arithmetic
- If you never used R before, go through Appendix A “A Sample Session” of the R manual on your computer

How Gauss became famous?





- 1801, *Dr Carl Friedrich Gauss*, 24; proved Fundamental Theorem of Algebra; wrote the book *Disquisitiones Arithmeticae*, which is still being studied today
- 1801, Jan 1 - Feb 11 (41 days), astronomer Piazzi observed Ceres (a dwarf planet), which was then lost behind sun
- 1801, Aug – Sep, futile search by top astronomers; Laplace claimed it unsolvable
- 1801, Oct – Nov, Gauss did calculations by *method of least squares*
- 1801, Dec 31, astronomer von Zach relocated Ceres according to Gauss' calculation
- 1802, *Summarische Übersicht der Bestimmung der Bahnen der beiden neuen Hauptplaneten angewandten Methoden*, considered the origin of linear algebra
- 1807, Professor of Astronomy and Director of Göttingen Observatory in remainder of his life
- 1809, *Theoria motus corporum coelestium in sectionibus conicis solem ambientum* (Theory of motion of the celestial bodies moving in conic sections around the Sun); birth of the Gaussian (normal) distribution, as an attempt to rationalize the method of least squares
- 1810, Laplace consolidated importance of Gaussian distribution by proving the central limit theorem

- 1829, Gauss-Markov Theorem. Under Gaussian error assumption (actually only uncorrelated and homoscedastic needed), least square solution is the best linear unbiased estimate (BLUE), i.e., it has the smallest variance and thus MSE among all linear unbiased estimators. Note other estimators such as the James-Stein estimator may have smaller MSE, but they are *nonlinear*.

For more details

- <http://www.keplersdiscovery.com/Asteroid.html>
- Teets and Whitehead (1999)

ARTICLES

The Discovery of Ceres: How Gauss Became Famous

DONALD TEETS
KAREN WHITEHEAD
South Dakota School of Mines and Technology
Rapid City, SD 57701

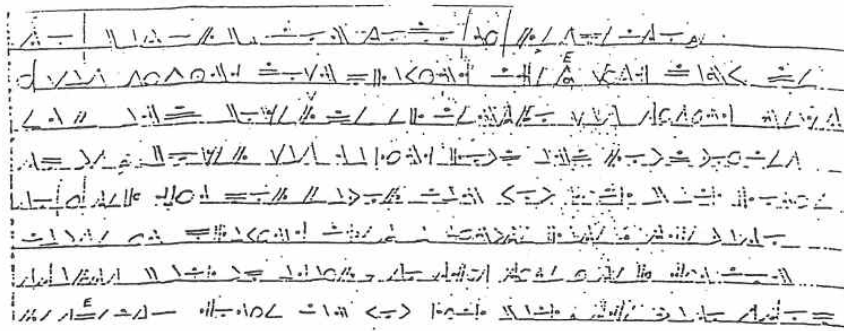
“The Duke of Brunswick has discovered more in his country than a planet: a super-terrestrial spirit in a human body.”

These words, attributed to Laplace in 1801, refer to the accomplishment of Carl Friedrich Gauss in computing the orbit of the newly discovered planetoid *Ceres Ferdinandea* from extremely limited data. Indeed, although Gauss had already achieved some fame among mathematicians, it was his work on the Ceres orbit that “made Gauss a European celebrity—this a consequence of the popular appeal which astronomy has always enjoyed . . .” [2]. The story of Gauss’s work on this problem is a good one and is often told in biographical sketches of Gauss (e.g., [2], [3], [6]), but the mathematical details of how he solved the problem are invariably omitted from such historical works. We are left to wonder: how did he do it? *Just how did Gauss*

Gauss’ story

- Motivated by a real problem.
- Heuristic solution: method of least squares.
- Solution readily verifiable: Ceres was re-discovered!
- *Algorithmic development*: linear algebra, Gaussian elimination, FFT (fast Fourier transform).
- Theoretical justification: Gaussian distribution, Gauss-Markov theorem.

A sampler by Marc Coram



ENTER HAMLET HAM TO BE OR NOT TO BE THAT IS THE QUESTION WHETHER TIS
 NOBLER IN THE MIND TO SUFFER THE SLINGS AND ARROWS OF OUTRAGEDUS
 FORTUNE OR TO TAKE ARMS AGAINST A SEA OF TROUBLES AND BY OPPOSING END

100 ER ENOHDLAE OHDL0 VOZEOUNORU 0 VOZEO HD OITO HEQQSET IUROFHE HENO ITORUZAEN
 200 ES ELOHRNDE OHRNO UOVEOULOSU 0 UOVEO HR OITO HEQQAET IUSOPHE HELO ITOSUVDEL
 300 ES ELOGRANDE OHANO UOVEOULOSU 0 UOVEO HA OITO HEQQAET IUSOPHE HELO ITOSUVDEL
 400 ES ELOHINME OHINO UOVEOULOSU 0 UOVEO HI OATO HEQQAET AUSOWHE HELO ATOSUVKEL
 500 ES ELOHINME OHINO UOVEOULOSU 0 UODEO HI OATO HEQQAET AUSOWHE HELO ATOSUDMEL
 600 ES ELOHINME OHINO UOVEOULOSU 0 UODEO HI OATO HEQQAET AUSOWHE HELO ATOSUDMEL
 900 ES ELOHANME OHANO UOVEOULOSU 0 UODEO HA OITO HEQQAET IUSOWHE HELO ITOSUDMEL
 1000 IS ILOHANMI OHANO RODIORLQSR 0 RODIO HA OETO HIOQUIT ERSOWHI HILÓ ETOSRDHIL
 1100 ISTILOHANMITOHANOT ODIO LOS TOT ODIOTHATOEROTHIOQUIRTE SOVHITHILOTEROS DMIL
 1200 ISTILOHANMITOHANOT ODIO LOS TOT ODIOTHATOEROTHIOQUIRTE SOVHITHILOTEROS DMIL
 1300 ISTILOHARMITOHANOT ODIO LOS TOT ODIOTHATOENOTHIOQUINTE SOVHITHILOTEROS DMIL
 1400 ISTILOHAMRITOHANOT OFIO LOS TOT OFIOTHATOENOTHIOQUINTE SOVHITHILOTEROS DMIL
 1600 ESTEL HAMRET HAM TO CE OL SOT TO CE THAT IN THE QUENTIOS WHEHTEL TIM SOCREL
 1700 ESTEL HAMRET HAM TO BE OL SOT TO BE THAT IN THE QUENTIOS WHEHTEL TIM SOCREL
 1800 ESTER HAMLET HAM TO BE OR SOT TO BE THAT IN THE QUENTIOS WHETHER TIM SOBREL
 1900 ENTER HAMLET HAM TO BE OR NOT TO BE THAT IS THE QUESTION WHETHER TIS NOBLER
 2000 ENTER HAMLET HAM TO BE OR NOT TO BE THAT IS THE QUESTION WHETHER TIS NOBLER

to bat-rb. con todo mi respeto. i was sitting down playing chess with
 danny de emf and boxer de el centro was sitting next to us. boxer was
 making loud and loud voices so i tell him por favor can you kick back
 homie cause im playing chess a minute later the vato starts back up again
 so this time i tell him con respecto homie can you kick back. the vato
 stop for a minute and he starts up again so i tell him check this out shut
 the f**k up cause im tired of your voice and if you got a problem with it
 we can go to celda and handle it. i really felt disrespected thats why i
 told him. anyways after i tell him that the next thing I know that vato
 slashes me and leaves. dy the time i figure im hit i try to get away but
 the c.o. is walking in my direction and he gets me right dy a celda. so i
 go to the hole. when im in the hole my home boys hit doxer so now "b" is
 also in the hole. while im in the hole im getting schoold wrong and

- Consulting project by Marc Coram (then a graduate student in statistics at Stanford); customer is a professor in political science.
- Marc modeled letter sequence by a Markov chain (26×26 transition matrix) and estimated transition probabilities from *War and Peace*.
- Now each mapping σ yields a likelihood $f(\sigma)$ of the symbol sequence.

- Find the σ that maximizes f . Sample space is at least $26! = 4.0329 \times 10^{26}$. Combinatorial optimization – hard!
- *Metropolis sampling*: At each iteration, generate a new σ' by random transposition of two letters; accept σ' with probability $\min \left\{ \frac{f(\sigma')}{f(\sigma)}, 1 \right\}$

Marc Coram’s story

- Motivated by a real problem.
- Solution readily verifiable: we can read it!
- *Algorithm*: Metropolis is one of top 10 algorithms in the 20th century.
- Read Diaconis (2009) for more details.

What is this course about?

- Not a course on “packages and languages for data analysis”. It does not answer questions such as “How to fit a linear mixed model in SAS or R?”
- Not a programming course, although it is extremely important and we do homework and projects in R.
- This course is about “numerical methods in statistics”. Our focus is on *algorithms*.

The form of a mathematical expression and the way the expression should be evaluated in actual practice may be quite different.

For a common numerical task in statistics, say solving the normal equation $\mathbf{X}^T \mathbf{X} \boldsymbol{\beta} = \mathbf{X}^T \mathbf{y}$, we need to know which methods/algorithms are there and what are their advantages and disadvantages. You will automatically fail this course if you use

```
solve(t(X) %*% X) %*% t(X) %*% y
```

Science, 287 # 5454, Feb 4, 2000, p799

Algorithms for the Ages

"Great algorithms are the poetry of computation," says Francis Sullivan of the Institute for Defense Analyses' Center for Computing Sciences in Bowie, Maryland. He and Jack Dongarra of the University of Tennessee and Oak Ridge National Laboratory have put together a sampling that might have made Robert Frost beam with pride—had the poet been a computer jock. Their list of 10 algorithms having "the greatest influence on the development and practice of science and engineering in the 20th century" appears in the January/February issue of *Computing in Science & Engineering*. If you use a computer, some of these algorithms are no doubt crunching your data as you read this. The drum roll, please:

- 1946: The Metropolis Algorithm for Monte Carlo. Through the use of random processes, this algorithm offers an efficient way to stumble toward answers to problems that are too complicated to solve exactly.
- 1947: Simplex Method for Linear Programming. An elegant solution to a common problem in planning and decision-making.
- 1950: Krylov Subspace Iteration Method. A technique for rapidly solving the linear equations that abound in scientific computation.
- 1951: The Decompositional Approach to Matrix Computations. A suite of techniques for numerical linear algebra.
- 1957: The Fortran Optimizing Compiler. Turns high-level code into efficient computer-readable code.
- 1959: QR Algorithm for Computing Eigenvalues. Another crucial matrix operation made swift and practical.
- 1962: Quicksort Algorithms for Sorting. For the efficient handling of large databases.
- 1965: Fast Fourier Transform. Perhaps the most ubiquitous algorithm in use today, it breaks down waveforms (like sound) into periodic components.
- 1977: Integer Relation Detection. A fast method for spotting simple equations satisfied by collections of seemingly unrelated numbers.
- 1987: Fast Multipole Method. A breakthrough in dealing with the complexity of n -body calculations, applied in problems ranging from celestial mechanics to protein folding.

Syllabus

Check course website frequently for updates and announcements.

<http://hua-zhou.github.io/teaching/st758-2014fall/>

Lecture notes will be updated and posted after each lecture.

2 Lecture 2, Sep 2

Announcements

- TA office hours, Wed @ 10A-12P?
- HW1 posted; due Sep 11 in class.

Last time

- Introduction. Gauss (least squares to find Ceres) – optimization, Marc Coram (decipher a note circulating in jail) – sampling. Two major modes of statistical computing.
- Course content, logistics.

Today

- Computer representation of characters, integers, and real numbers.

Computer storage and arithmetic

Elementary units of computer storage: *bit* = “binary” + “digit” (coined by John Tukey), *byte* = 8 bits, kB = kilobyte = 10^3 bytes, MB = megabytes = 10^6 bytes, GB = gigabytes = 10^9 bytes, TB = terabytes = 10^{12} bytes, PB = petabytes = 10^{15} bytes, ...

Storage of characters

ASCII control characters				ASCII printable characters				Extended ASCII characters									
00	NULL	(Null character)		32	space	64	@	96	`	128	Ç	160	á	192	Ł	224	Ó
01	SOH	(Start of Header)		33	!	65	A	97	a	129	ú	161	í	193	ł	225	ó
02	STX	(Start of Text)		34	"	66	B	98	b	130	é	162	ó	194	Ł	226	Ô
03	ETX	(End of Text)		35	#	67	C	99	c	131	â	163	ú	195	ł	227	Ö
04	EOT	(End of Trans.)		36	\$	68	D	100	d	132	ä	164	ñ	196	—	228	ö
05	ENQ	(Enquiry)		37	%	69	E	101	e	133	à	165	Ñ	197	†	229	Ø
06	ACK	(Acknowledgement)		38	&	70	F	102	f	134	á	166	ª	198	‡	230	µ
07	BEL	(Bell)		39	'	71	G	103	g	135	ç	167	º	199	Ä	231	þ
08	BS	(Backspace)		40	(72	H	104	h	136	ê	168	¿	200	Å	232	þ
09	HT	(Horizontal Tab)		41)	73	I	105	i	137	ë	169	®	201	Æ	233	Û
10	LF	(Line feed)		42	*	74	J	106	j	138	è	170	™	202	Œ	234	Ü
11	VT	(Vertical Tab)		43	+	75	K	107	k	139	ï	171	¼	203	Œ	235	Ù
12	FF	(Form feed)		44	,	76	L	108	l	140	î	172	½	204	Œ	236	Ý
13	CR	(Carriage return)		45	-	77	M	109	m	141	í	173	¾	205	Œ	237	ÿ
14	SO	(Shift Out)		46	.	78	N	110	n	142	Ä	174	€	206	Œ	238	—
15	SI	(Shift In)		47	/	79	O	111	o	143	Å	175	€	207	Œ	239	—
16	DLE	(Data link escape)		48	0	80	P	112	p	144	É	176	€	208	Œ	240	≡
17	DC1	(Device control 1)		49	1	81	Q	113	q	145	æ	177	€	209	Œ	241	±
18	DC2	(Device control 2)		50	2	82	R	114	r	146	Æ	178	€	210	Œ	242	—
19	DC3	(Device control 3)		51	3	83	S	115	s	147	ó	179	€	211	Œ	243	¾
20	DC4	(Device control 4)		52	4	84	T	116	t	148	ô	180	€	212	Œ	244	¾
21	NAK	(Negative acknowl.)		53	5	85	U	117	u	149	õ	181	€	213	Œ	245	§
22	SYN	(Synchronous idle)		54	6	86	V	118	v	150	ù	182	€	214	Œ	246	+
23	ETB	(End of trans. block)		55	7	87	W	119	w	151	ú	183	€	215	Œ	247	—
24	CAN	(Cancel)		56	8	88	X	120	x	152	ÿ	184	€	216	Œ	248	—
25	EM	(End of medium)		57	9	89	Y	121	y	153	Œ	185	€	217	Œ	249	—
26	SUB	(Substitute)		58	:	90	Z	122	z	154	Œ	186	€	218	Œ	250	—
27	ESC	(Escape)		59	;	91	[123	{	155	ø	187	€	219	Œ	251	—
28	FS	(File separator)		60	<	92	\	124		156	€	188	€	220	Œ	252	—
29	GS	(Group separator)		61	=	93]	125	}	157	Œ	189	€	221	Œ	253	—
30	RS	(Record separator)		62	>	94	^	126	~	158	×	190	€	222	Œ	254	—
31	US	(Unit separator)		63	?	95	_			159	f	191	€	223	Œ	255	nbsp
127	DEL	(Delete)															

- Plain text files are stored in the form of characters: `.r`, `.c`, `.cpp`, `.tex`, `.html`, ...
- ASCII (American Code for Information Interchange): 7 bits, only $2^7 = 128$ characters, “Hua” corresponds to “48 75 61 (Hex) = 72 117 97 (Dec) = 1001000 1110101 1100001”.
- Extended ASCII: 8 bits, $2^8 = 256$ characters.
- Unicode: UTF-8, UTF-16 and UTF-32 support many more characters including foreign characters; last 7 digits conform to ASCII. UTF-8 is the current dominant character encoding on internet.

Fixed-point number system

Fixed-point number system \mathbb{I} is a computer model for integers \mathbb{Z} . One storage unit may be $M = 8/16/32/64$ bit.

- The number of bits and method of representing negative numbers vary from system to system. MATLAB has `(u)int8`, `(u)int16`, `(u)int32`, `(u)int64`. The `integer` type in R has $M = 32$ bits.

- First bit indicates sign: 0 for nonnegative numbers, 1 for negative numbers.
- “two’s complement representation” for negative numbers. (i) Sign bit is set to 1, (ii) remaining bits are set to opposite values, (iii) 1 is added to the result.

+18		0 0 0 1 0 0 1 0 <div style="display: flex; justify-content: center; gap: 20px;"> ↑ Sign bit is 0 Binary equivalent of +18 </div>
-18	Signed magnitude representation	1 0 0 1 0 0 1 0 <div style="display: flex; justify-content: center; gap: 20px;"> ↑ Sign bit is 1 Binary equivalent of +18 </div>
	Signed 1's complement representation	1 1 1 0 1 1 0 1 <div style="display: flex; justify-content: center; gap: 20px;"> ↑ Sign bit is 1 1's complement of +18 </div>
	Signed 2's complement representation	1 1 1 0 1 1 1 0 <div style="display: flex; justify-content: center; gap: 20px;"> ↑ Sign bit is 1 2's complement of +18 </div>

- Range of representable integers by M -bit storage unit is $[-2^{M-1}, 2^{M-1} - 1]$ (don't need to represent 0 anymore so could have capacity for 2^{M-1} negative numbers).
- For $M = 8$, $[-128, 127]$. For $M = 16$, $[-65536, 65535]$.
For $M = 32$, $[-2147483648, 2147483647]$.
- Following code snippet shows that the smallest integer in R is $-2^{31} + 1 = -2147483647$.

```
# .Machine variable  
.Machine$integer.max
```

```
[1] 2147483647
```

```
# integer type in R uses M=32 bits  
M <- 32  
big <- 2^(M-1) - 1  
small <- - 2^(M-1)  
as.integer(big)
```

```
[1] 2147483647
```

```
as.integer(big+1)
```

```
Warning: NAs introduced by coercion
```

```
[1] NA
```

```
as.integer(small+1)
```

```
[1] -2147483647
```


```
as.integer(small)
```

```
Warning: NAs introduced by coercion
```

```
[1] NA
```

- For unsigned integers such as in MATLAB, the range is $[0, 2^M - 1]$.

Two's Complement representation using 4 bit binary strings




- An integer $-i$ in the interval $[-2^{M-1}, -1]$ would be represented by the same bit pattern by which the nonnegative integer $2^M - i$ is represented, treating the sign bit as a regular numeric bit. For example, $(-18) = 11101110$, $(2^8) - (18) = (238) = 11101110$. Two's complement is subtracting the nonnegative integer i from $1 \dots 1 + 1 = 10 \dots 0 = (2^M)$.
- Addition and subtraction are much simpler in two's complement representation. Why? it respects modular arithmetic nicely (look at the diagram). E.g., $(3) + (4) = 0011 + 0100 = 0111 = (7)$. Need to keep track of crossing two boundaries $10 \dots 0$ and $00 \dots 0$. Cross boundary $10 \dots 0$ once, we add 2^M , e.g., $(6) + (7) = 0110 + 0111 = 1101 = (-3) + (2^4) = (13)$. Cross boundary $00 \dots 0$, we subtract

2^M , e.g., $(-6) + (-7) = 1010 + 1001 = 0011 = (3) - (2^4) = (-13)$. Crossing both boundaries is fine, e.g., $(-3)+(-4) = 1101 + 1100 = 1001 = (-7)$.

- Keep track of overflow and underflow. If the result of a summation is R , which must be in the set $[-2^{M-1}, 2^{M-1} + 1]$, there are only three possibilities for the true sum: R , $R + 2^M$ (overflow), or $R - 2^M$ (underflow).
 - When adding two nonnegative integers: $0XX\dots X + 0YY\dots Y$, where X and Y are arbitrary binary digits, we only need to keep track of overflow in this case. If the resulting binary number has a leading bit of 1, we know overflow occurs since the sum cannot be negative. We should treat that sign bit as a regular numeric bit. In other words, we crossed the upper boundary 100 and should add 2^M to the result. If the resulting binary number has a leading bit of 0, no overflow occurs.
 - When adding two negative integers: $1XX\dots X + 1YY\dots Y$, where X and Y are arbitrary binary digits, we only need to keep track of underflow in this case. If the resulting binary number has a leading bit of 0, we know underflow occurs since the sum cannot be positive. We crossed the lower boundary 000 and should subtract 2^M from the result. If the resulting binary number has a leading bit of 1, no underflow occurs.
 - When adding a negative integer and a nonnegative integer: $1XX\dots X + 0YY\dots Y$, the result is always between the two summands. No overflow or underflow could happen.

Floating-point number system



Floating-point number system \mathbb{F} is a computer model for real numbers \mathbb{R} .

- A real is represented by $\pm d_0.d_1d_2 \cdots d_p \times b^e$ (scientific notation).
- Parameters for a floating-point number system: *base* (or *radix*), range of *fraction* (or *mantissa*, *significand*), range of *exponent*.
- Non-uniqueness of representation. *normalized/denormalized* significant digits. E.g., $+18 = +1.0010 \times 2^4$ (normalized) $= +0.10010 \times 2^5$ (denormalized).
- *Bias* (or *excess*): actual exponent is obtained by subtracting bias from the value of exponent evaluated regardless of sign digit.
- IEEE 754-1985 and IEEE 754-2008.
 - *Single precision* (32 bit): base 2, $p = 23$ (23 significant bits), $e_{\max} = 127$, $e_{\min} = -126$ (8 exponent bits), bias=127. $e_{\min} - 1$ and $e_{\max} + 1$ are reserved for special numbers. This implies a maximum magnitude of $\log_{10}(2^{127}) \approx 38$ and precision to $\log_{10}(2^{23}) \approx 7$ decimal point. $\pm 10^{\pm 38}$.
 - *Double precision* (64 bit): base 2, $p = 52$ (52 significant bits), $e_{\max} = 1023$, $e_{\min} = -1022$ (11 exponent bits), bias=1023. This implies a maximum magnitude of $\log_{10}(2^{1023}) \approx 308$ and precision to $\log_{10}(2^{52}) \approx 16$ decimal point. $\pm 10^{\pm 308}$.

- “(+18) = (2⁴ + 2¹) = +1.0010 × 2⁴ in single precision

$$(0)(10000011)(001000000000000000000000).$$

First is sign bit. Next 8 bits are exponent 131 in ordinary base 2 with a bias of 127. Remaining 23 bits represent the fraction beyond the leading bit, known to be 1. In summary it represents (+18) as +1.0010 × 2⁴ in the binary format. (−18) is represented by the same bits except changing the sign bit to 1.

- Special floating-point numbers:
 - Exponent $e_{\max} + 1$ plus a mantissa of 0 means $\pm\infty$.
 - Exponent $e_{\max} + 1$ plus a nonzero mantissa means NaN. NaN could be produced from $0/0$, $0*\text{Inf}$, ... In general $\text{NaN} \neq \text{NaN}$.
 - Exponent $e_{\min} - 1$ with a mantissa of all 0s represents the real number 0.
 - Exponent $e_{\min} - 1$ with a nonzero mantissa are for numbers less than $b^{e_{\min}}$. Numbers are de-normalized in the range $(0, b^{e_{\min}})$ – “graceful underflow”.
- \mathbb{F} is not a subset of \mathbb{R} , although $\mathbb{I} \subset \mathbb{Z}$.
- R only uses double (64-bit) and 32-bit integer. It can be a downside when dealing with big data.

To summarize

- Single precision: $\pm 10^{\pm 38}$ with precision up to 7 decimal digits.
- Double precision: $\pm 10^{\pm 308}$ with precision up to 16 decimal digits.
- Irrational numbers such as π do not exist in \mathbb{F} .
- Exercise: what is the floating point representation of the number 0.1?
- The floating-point numbers do not occur uniformly over the real number line.




Fig. 2.4. The Floating-Point Number Line, Nonnegative Half




Fig. 2.5. The Floating-Point Number Line, Nonpositive Half




Fig. 2.6. The Floating-Point Number Line, Nonnegative Half; Another View




Fig. 2.7. The Fixed-Point Number Line, Nonnegative Half

- *Machine epsilons* are the spacings of numbers around 1. $\epsilon_{\min} = b^{-p}$ and $\epsilon_{\max} = b^{1-p}$.




Fig. 2.8. Relative Spacings at 1: “Machine Epsilons”

- The variable `.Machine` in `R` contains numerical characteristics of the machine.
- How to test `inf` and `nan`? In `R`, `is.nan()`, `is.finite()`, `is.infinite()`. In `MATLAB`, `isinf()`, `isnan()`.

3 Lecture 3, Sep 4

Announcements

- TA office hours: Wed @ 1P-3P or Wed @ 10A-12P?

Last time

- Computer representation of characters
- Fixed-point number system for integers
- Floating-point number system (IEEE 754 standards for single/double precision)

Today


- Consequences of compute storage/arithmetic
- Algorithms
- Computer languages

Consequences of computer storage/arithmetic

- Be memory conscious when dealing with big data. E.g., human genome has about 3×10^9 bases, each of which belongs to $\{A, C, T, G\}$. How much storage if we store 10^6 SNPs (single nucleotide polymorphisms) of 1000 individuals (1000 Genome Project) as characters (1GB), single (4GB), double (8GB), int32 (4GB), int16 (2GB), int8 (1GB), PLINK binary format 2bit/SNP (250MB)?
- Know the limit. *Overflow* and *Underflow*. For double precision, $\pm 10^{\pm 308}$. In most situations, underflow is preferred over overflow. Overflow often causes crashes. Underflow yields zeros. E.g., in logistic regression, $p_i = \frac{\exp(\mathbf{x}_i^T \boldsymbol{\beta})}{1 + \exp(\mathbf{x}_i^T \boldsymbol{\beta})} = \frac{1}{1 + \exp(-\mathbf{x}_i^T \boldsymbol{\beta})}$. The former expression can easily lead to $\infty/\infty = NaN$, while the latter expression leads to graceful underflow.
- Be aware of non-uniform distribution of floating-point numbers, in contrast to fixed-point numbers. There are the same number of floating-point numbers in

$[b^i, b^{i+1}]$ and $[b^{i+1}, b^{i+2}]$ for $e_{\min} \leq i \leq e_{\max} - 2$. It is more dense when closer to zero.

- “Catastrophic cancellation 1”. Addition or subtraction of two numbers of widely different magnitudes: $a + b$ or $a - b$ where $a \gg b$ or $a \ll b$. We lose the precision in the number of smaller magnitude. Consider $a = x.xxxx\dots \times 2^0$ and $b = y.yyy\dots \times 2^{-53}$. What happens when computer calculates $a + b$? We get $a + b = a$!
- Another example: What happens when compute $\sum_{x=1}^{\infty} x$ in order? Will the partial sum reach Inf? “A divergent series converges.”
- Always try to add numbers of similar magnitude. Rule 1: add small numbers together before adding larger ones. Rule 2: add numbers of like magnitude together (paring). When all numbers are of same sign and similar magnitude, add in pairs so each stage the summands are of similar magnitude.



- “Catastrophic cancellation 2”. Subtraction of two nearly equal numbers eliminates significant digits. $a - b$ where $a \approx b$. Consider $a = x.xxxxxxxxxxxx1ssss$, $b = x.xxxxxxxxxxxx0tttt$. The result is $1.vvvvu\dots u$ where u are unassigned digits.
- Violation of associative law and distributive law.

More about cancellation

- Evaluating e^{-20} by Taylor series

$$e^{-x} = 1 - x + x^2/2! - x^3/3! + \dots$$

gets **6.138e-09**, while the true value is about **2.061e-09**. Many cancellations accumulate. Anyway, it's *not* the way to evaluate e^x !

- Sometimes catastrophic cancellation can be avoided. Roots of the quadratic function $ax^2 + bx + c$ are

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}.$$

When one root is close to 0, cancellation can happen. We may evaluate one of the root (away from 0) by the formula and then compute the other by relationship $x_1x_2 = c/a$.

Algorithms

- *Algorithm* is loosely defined as a set of instructions for doing something. Input \rightarrow Output.
- Knuth (2005): (1) finiteness, (2) definiteness, (3) input, (4) output, (5) effectiveness
- Basic unit for measuring efficiency is flop. A *flop* (floating point operation) consists of a floating point multiply (or divide) and the usually accompanying addition, fetch and store. Some books such as Lange (2010) and Golub and Van Loan (2013) consider addition as a separate flop.
- How to measure efficiency of an algorithm? Big O notation. If n is the size of a problem, an algorithm has order $O(f(n))$ if, as $n \rightarrow \infty$, the number of computations $\rightarrow cf(n)$, where c is some constant that does not depend on n .
- E.g., matrix-vector multiplication $\mathbf{A}*\%*\mathbf{b}$, where $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^n$, takes $O(mn)$ flops. Matrix-matrix multiplication $\mathbf{A}*\%*\mathbf{B}$, where $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{B} \in \mathbb{R}^{n \times p}$, takes $O(mnp)$ flops.
- *Exponential* order $O(b^n)$ (NP-hard=“horrible”), *polynomial* order $O(n^q)$ (doable), $O(n \log n)$ (fast), linear order $O(n)$ (fast), *log* order $O(\log n)$ (fast).
- One goal of this course is to get familiar with the flop counts for common numerical tasks in statistics.

The form of a mathematical expression and the way the expression should be evaluated in actual practice may be quite different.

- Compare flops of the following two expressions:

$\mathbf{G} \mathbf{Xt} \mathbf{y}$

$\mathbf{G} (\mathbf{Xt} \mathbf{y})$

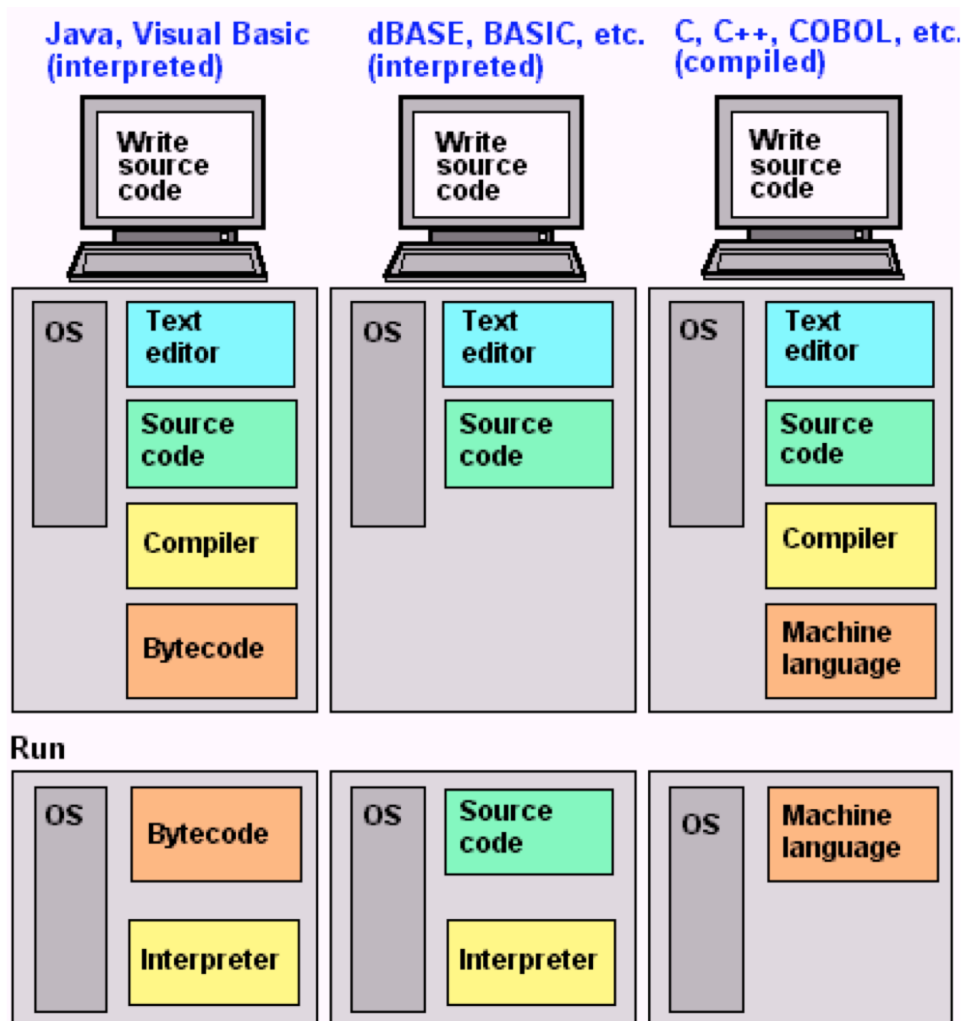
where $\mathbf{G} \in \mathbb{R}^{p \times p}$, $\mathbf{Xt} \in \mathbb{R}^{p \times n}$, and $\mathbf{y} \in \mathbb{R}^n$. “Matrix multiplication is expensive.”

- Hardware advancement, e.g., CPU clock rate, only affects constant c . Unfortunately, data size n is increasing too and often at a faster rate.
- Classification of data sets by Huber (1994, 1996).

Data Size	Bytes	Storage Mode
Tiny	10^2	Piece of paper
Small	10^4	A few pieces of paper
Medium	10^6 (megabyte)	A floppy disk
Large	10^8	Hard disk
Huge	10^9 (gigabytes)	Hard disk(s)
Massive	10^{12} (terabytes)	RAID storage

- It is an era of “big data”: wiki, WSJ, white house, McKinsey report, ..., meaning great opportunities for statisticians. However we should be aware of the gap between classical statistics curriculum and reality.
- Difference of $O(n^2)$ and $O(n \log n)$ on massive data. Suppose we have a teraflop supercomputer – capable of doing 10^{12} flops per second. For a problem of size $n = 10^{12}$, $O(n \log n)$ algorithm takes about $10^{12} \log(10^{12})/10^{12} \approx 27$ seconds. $O(n^2)$ algorithm takes about 10^{12} seconds, which is approximately 31710 years!
- QuickSort and FFT are celebrated algorithms that turn $O(n^2)$ operations into $O(n \log n)$. “Divide-and-conquer” is a powerful technique. Another example is the Strassen’s method, which turns $O(n^3)$ matrix multiplication into $O(n^{\log_2 7})$.

Computer languages



Compiled versus interpreted languages.

- Compiled languages: C/C++, Fortran, ... Directly compiled to machine code that is executed by CPU. Advantage: fast, take less memory. Disadvantage: relatively longer development time, hard to debug.
- Interpreted language: R, Matlab, SAS IML, ... Interpreted by interpreter. Advantage: fast for prototyping. Disadvantage: excruciatingly slow for loops.
- Mixed (compiled and then interpreted by virtual machine): Python, JAVA. Advantage: extremely convenient for data preprocessing and manipulation; relatively short development time. Disadvantage: not as fast as compiled language.

- Scripting: Unix/Linux scripts, Perl, Python. Extremely useful for data pre-processing and manipulation. E.g., massage the Yelp (http://www.yelp.com/dataset_challenge) data before analysis

```

hzhou3@zhou-lap2:~/yelp_phoenix_academic_dataset $ ls -all
total 452432
drwx-----@ 9 hzhou3  staff      306 Mar 28 10:55 ./
drwxr-xr-x  7 hzhou3  staff      238 Apr 10 09:10 ../
-rw-r--r--@ 1 hzhou3  staff      6148 Mar 28 10:55 .DS_Store
-rw-r--r--@ 1 hzhou3  staff     140579 Mar 26 18:44 READ_FIRST-Phoenix_Academic_Dataset_Agreement-3-11-13.pdf
-rw-r--r--  1 hzhou3  staff       421 Mar 26 17:10 notes.txt
-rw-r--r--@ 1 hzhou3  staff    4287324 Mar 28 13:19 yelp_academic_dataset_business.json
-rw-r--r--  1 hzhou3  staff    3519126 Mar 26 17:54 yelp_academic_dataset_checkin.json
-rw-r--r--@ 1 hzhou3  staff   216292386 Mar 28 13:30 yelp_academic_dataset_review.json
-rw-r--r--  1 hzhou3  staff    7374045 Mar 26 17:54 yelp_academic_dataset_user.json
hzhou3@zhou-lap2:~/yelp_phoenix_academic_dataset $ head yelp_academic_dataset_user.json
{"votes": {"funny": 0, "useful": 7, "cool": 0}, "user_id": "CR2y7yEm4X035ZMrTtN9Q", "name": "Jim", "average_stars": 5.0, "review_count": 6, "type": "user"}
{"votes": {"funny": 0, "useful": 1, "cool": 0}, "user_id": "_9GXoHhdx30ujPaQwh6Ew", "name": "Kelle", "average_stars": 1.0, "review_count": 2, "type": "user"}
{"votes": {"funny": 0, "useful": 1, "cool": 0}, "user_id": "8mM-nqxjg6pT04kwcjMbsw", "name": "Stephanie", "average_stars": 5.0, "review_count": 2, "type": "user"}
{"votes": {"funny": 0, "useful": 2, "cool": 0}, "user_id": "Ch6CdTR2IVaVANr-RgLM0g", "name": "T", "average_stars": 5.0, "review_count": 2, "type": "user"}
{"votes": {"funny": 0, "useful": 0, "cool": 0}, "user_id": "NZrLmHRYiHmyT1JrfzkCOA", "name": "Beth", "average_stars": 1.0, "review_count": 1, "type": "user"}
{"votes": {"funny": 30, "useful": 45, "cool": 36}, "user_id": "mWx5Sxt_dx-sYBZg6RgJHQ", "name": "Amy", "average_stars": 3.79, "review_count": 19, "type": "user"}
{"votes": {"funny": 28, "useful": 130, "cool": 31}, "user_id": "hryUDaRk7FLuDAYui2oldw", "name": "Beach", "average_stars": 3.8300000000000001, "review_count": 207, "type": "user"}
{"votes": {"funny": 1, "useful": 0, "cool": 1}, "user_id": "2t6fZNLtiqsihVme07zggg", "name": "christine", "average_stars": 3.0, "review_count": 2, "type": "user"}
{"votes": {"funny": 0, "useful": 3, "cool": 2}, "user_id": "mn6F-eP5WU37b-iLTop2mQ", "name": "Denis", "average_stars": 4.5, "review_count": 4, "type": "user"}
{"votes": {"funny": 5, "useful": 24, "cool": 9}, "user_id": "myXq7PFXkD_yfXT580SXmW", "name": "Shawn", "average_stars": 3.8999999999999999, "review_count": 10, "type": "user"}
hzhou3@zhou-lap2:~/yelp_phoenix_academic_dataset $

```

- Database language: SQL, Hadoop. Data analysis never happens if we do not know how to retrieve data from databases.

4 Lecture 4, Sep 9

Announcements

- TA office hours. Fri @ 10A-12P, Wed @ 1P-3P, Wed @ 10A-12P, or Thu @ 1P-2P?
- Some resources for R programming:
 - Dr. John Monahan's class (2013 fall) on R: <http://www.stat.ncsu.edu/people/monahan/courses/ST610/>
 - Code school: <http://tryr.codeschool.com/levels/1/challenges/1>
 - *Advanced R* by Hadley Wickham: <http://adv-r.had.co.nz/>

Last time

- Consequences of computer storage/arithmetics (overflow, underflow, catastrophic cancellation, ...)
- Algorithms, flops, big O notation
- Computer languages (compiled, interpreted, mixed, scripting)

Today

- Review of linear algebra (self-study and do HW2)
- Numerical linear algebra: preliminaries

More about computer languages

- To improve efficiency of interpreted languages such as R or MATLAB code, avoid loops as much as possible. Aka, vectorize code.
- For some tasks where looping is necessary (cannot vectorize code), consider coding in C/C++ or FORTRAN. It is convenient to incorporate compiled code into R or MATLAB.

- To be versatile in dealing with big data, master at least one language in each category. Take advantage of the resource in the department. E.g., check out Chris’s class on PYTHON, John’s class on R programming, ...
- Don’t reinvent wheels. Make good use of libraries BLAS, LAPACK/LINPACK, BOOST, SCIPY, NUMPY, ...
- One example of my recent pedigree GWAS project, which fits large variance components model on $n \sim 10^3$ persons and $\sim 10^6$ SNPs. MENDEL implements the method using FORTRAN, takes advantage of problem structure, and is flop frugal. GWAF (an R package) uses the canned routine in R for fitting variance component, which does not scale up.

# SNPs	MENDEL	GWAF-GEE	GWAF-LMM
100	4.81	0.71	8.83
1,000	4.80	7.71	87.06
10,000	5.40	207.60	894.82
100,000	10.78	26,486.92	11,703.88

Table 1: Comparison of run times (in seconds). Run times are based on testing the first 100, 1000, 10000 and 100000 SNPs on chromosome 19. Trait values are from the same simulation replicate in Figure 3.

Many aspects could contribute to this dramatic difference in efficiency: looping, compiled vs interpreted language, problem structure, choice of data types (remember R only does double precision), memory management,

- Distinction between compiled language and interpreted language is getting blurred. JIT (just-in-time) compilation technology in MATLAB since 2002 (v6.5). COMPILER package in R for JIT compiling since 2012.

Numerical linear algebra

- The first big chunk of this course is numerical linear algebra.
- Topics in numerical algebra: BLAS, solve linear equations $\mathbf{Ax} = \mathbf{b}$, regression computations $\mathbf{X}^T \mathbf{X} \boldsymbol{\beta} = \mathbf{X}^T \mathbf{y}$, eigen-problems $\mathbf{Ax} = \lambda \mathbf{x}$ and generalized eigen-problems $\mathbf{Ax} = \lambda \mathbf{Bx}$, singular value decompositions $\mathbf{A} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^T$, ...

- We start with review of some linear algebra facts (review by yourself and do HW2).

Vector and matrix norms

KL chapter 6. Norm measures the “size”.

- Vector norm $\|\cdot\| : \mathbb{R}^n \rightarrow \mathbb{R}$. (a) $\|\mathbf{x}\| \geq 0$, (b) $\|\mathbf{x}\| = 0$ if and only if $\mathbf{x} = \mathbf{0}$, (c) $\|c\mathbf{x}\| = c\|\mathbf{x}\|$ (homogeneity), (d) $\|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|$ (triangle inequality).
- ℓ_p norm. $\|\mathbf{x}\|_p = (\sum_i |x_i|^p)^{1/p}$, $p \in [1, \infty]$. ℓ_1 is the Manhattan norm. ℓ_2 is the Euclidean norm. ℓ_∞ is the sup norm.
- For matrix norm $\|\cdot\| : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$, we further require (e) $\|\mathbf{AB}\| \leq \|\mathbf{A}\|\|\mathbf{B}\|$.
- Frobenius norm $\|\mathbf{A}\|_F = (\sum_i \sum_j a_{ij}^2)^{1/2}$. Property of (e) is checked by Cauchy-Schwartz inequality.
- Induced matrix norm (or operator norm): $\|\mathbf{A}\| = \sup_{\mathbf{x} \neq \mathbf{0}} \frac{\|\mathbf{Ax}\|}{\|\mathbf{x}\|} = \sup_{\|\mathbf{x}\|=1} \|\mathbf{Ax}\|$ for any fixed vector norm. To check property (e), let $\mathbf{y} = \mathbf{Bx}$, then $\|\mathbf{AB}\| = \sup_{\mathbf{x} \neq \mathbf{0}} \frac{\|\mathbf{Ay}\|}{\|\mathbf{y}\|} \frac{\|\mathbf{Bx}\|}{\|\mathbf{x}\|} \leq \|\mathbf{A}\| \sup_{\mathbf{x} \neq \mathbf{0}} \frac{\|\mathbf{Bx}\|}{\|\mathbf{x}\|} = \|\mathbf{A}\|\|\mathbf{B}\|$.
- Matrix-1 norm, $\|\mathbf{A}\|_1 = \max_j \sum_i |a_{ij}|$.
Matrix-2 norm, $\|\mathbf{A}\|_2 = \sqrt{\rho(\mathbf{A}^T \mathbf{A})} = \max_{\|\mathbf{u}\|_2=1, \|\mathbf{v}\|_2=1} \mathbf{u}^T \mathbf{A} \mathbf{v}$, which reduces to $\rho(\mathbf{A})$ if \mathbf{A} is symmetric. ρ is the spectral radius of a matrix, the absolute value of the dominant eigenvalue.
Matrix- ∞ norm, $\|\mathbf{A}\|_\infty = \max_i \sum_j |a_{ij}|$.
- When \mathbf{A} is a column vector, these matrix induced norms reduce to the original vector norm.
- $\rho(\mathbf{A}) \leq \|\mathbf{A}\|$ for any induced matrix norm. For any \mathbf{A} and $\epsilon > 0$, there exists an induced matrix norm such that $\|\mathbf{A}\| \leq \rho(\mathbf{A}) + \epsilon$.

Rank

Assume $\mathbf{A} \in \mathbb{R}^{m \times n}$.

- $\text{rank}(\mathbf{A})$ is the maximum number of linearly independent rows (or columns) of a matrix.

- $\text{rank}(\mathbf{A}) \leq \min\{m, n\}$.
- A matrix is *full rank* if $\text{rank}(\mathbf{A}) = \min\{m, n\}$. It is *full row rank* if $\text{rank}(\mathbf{A}) = m$. It is *full column rank* if $\text{rank}(\mathbf{A}) = n$.
- A square matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ is *singular* if $\text{rank}(\mathbf{A}) < n$ and *non-singular* if $\text{rank}(\mathbf{A}) = n$.
- $\text{rank}(\mathbf{AB}) \leq \min\{\text{rank}(\mathbf{A}), \text{rank}(\mathbf{B})\}$. “Matrix multiplication cannot increase the rank.”
- $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{A}^T) = \text{rank}(\mathbf{A}^T \mathbf{A}) = \text{rank}(\mathbf{A} \mathbf{A}^T)$.
- $\text{rank}(\mathbf{AB}) = \text{rank}(\mathbf{A})$ if \mathbf{B} has full row rank.
- $\text{rank}(\mathbf{AB}) = \text{rank}(\mathbf{B})$ if \mathbf{A} has full column rank.
- $\text{rank}(\mathbf{A} + \mathbf{B}) \leq \text{rank}(\mathbf{A}) + \text{rank}(\mathbf{B})$.

Trace

$\mathbf{A} \in \mathbb{R}^{n \times n}$ a square matrix.

- $\text{tr}(\mathbf{A}) = \sum_{i=1}^n a_{ii}$
- $\text{tr}(\mathbf{A} + \mathbf{B}) = \text{tr}(\mathbf{A}) + \text{tr}(\mathbf{B})$
- $\text{tr}(\lambda \mathbf{A}) = \lambda \text{tr}(\mathbf{A})$ where λ is a scalar
- $\text{tr}(\mathbf{A}^T) = \text{tr}(\mathbf{A})$
- Invariance under cycle permutation: $\text{tr}(\mathbf{AB}) = \text{tr}(\mathbf{BA})$. In general, $\text{tr}(\mathbf{A}_1 \cdots \mathbf{A}_k) = \text{tr}(\mathbf{A}_{j+1} \cdots \mathbf{A}_k \mathbf{A}_1 \cdots \mathbf{A}_j)$.

Orthogonality and orthogonalization

- \mathbf{v}_1 is *orthogonal* to \mathbf{v}_2 , written $\mathbf{v}_1 \perp \mathbf{v}_2$, if $\langle \mathbf{v}_1, \mathbf{v}_2 \rangle = \mathbf{v}_1^T \mathbf{v}_2 = 0$. They are *orthonormal* if $\mathbf{v}_1 \perp \mathbf{v}_2$ and $\|\mathbf{v}_i\|_2 = 1$, $i = 1, 2$.

- Gram-Schmidt transformation orthonormalizes two non-zero vectors \mathbf{x}_1 and \mathbf{x}_2 .

$$\tilde{\mathbf{x}}_1 = \frac{1}{\|\mathbf{x}_1\|_2} \mathbf{x}_1$$

$$\tilde{\mathbf{x}}_2 = \frac{1}{\|\mathbf{x}_2 - \langle \tilde{\mathbf{x}}_1, \mathbf{x}_2 \rangle \tilde{\mathbf{x}}_1\|_2} (\mathbf{x}_2 - \langle \tilde{\mathbf{x}}_1, \mathbf{x}_2 \rangle \tilde{\mathbf{x}}_1)$$




Fig. 2.2. Orthogonalization of \mathbf{x}_1 and \mathbf{x}_2

- A set of nonzero, mutually orthogonal vectors are linearly independent.
- A real square matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ is orthogonal if $\mathbf{A}^T \mathbf{A} = \mathbf{I}_n$, i.e., its rows/columns are orthonormal. Orthogonal matrix is of full rank, thus $\mathbf{A}^T = \mathbf{A}^{-1}$ and $\mathbf{A} \mathbf{A}^T = \mathbf{I}_n$.

Positive (semi)definite matrix

Assume $\mathbf{A} \in \mathbb{R}^{n \times n}$ is *symmetric*.

- A real symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ is *positive semidefinite* (or *nonnegative definite*) if $\mathbf{x}^T \mathbf{A} \mathbf{x} \geq 0$ for all \mathbf{x} . Notation: $\mathbf{A} \succeq \mathbf{0}_{n \times n}$.
- E.g., the *Gramian matrix* $\mathbf{X}^T \mathbf{X}$ or $\mathbf{X} \mathbf{X}^T$.
- If inequality is strict for all $\mathbf{x} \neq \mathbf{0}$, then \mathbf{A} is *positive definite*. Notation: $\mathbf{A} \succ \mathbf{0}_{n \times n}$.
- $\mathbf{A} \succeq \mathbf{B}$ means $\mathbf{A} - \mathbf{B} \succeq \mathbf{0}_{n \times n}$.

- If $\mathbf{A} \succeq \mathbf{B}$, then $\det(\mathbf{A}) \geq \det(\mathbf{B})$ with equality if and only if $\mathbf{A} = \mathbf{B}$.
- $\mathbf{A} \in \mathbb{R}^{n \times n}$ is positive semidefinite if and only if \mathbf{A} is a covariance matrix of a random vector in \mathbb{R}^n .

Matrix inverses

Assume $\mathbf{A} \in \mathbb{R}^{m \times n}$.

- The *Moore-Penrose inverse* of \mathbf{A} is a matrix $\mathbf{A}^+ \in \mathbb{R}^{n \times m}$ with following properties
 - (a) $\mathbf{A}\mathbf{A}^+\mathbf{A} = \mathbf{A}$. (*Generalized inverse, g_1 inverse, or inner pseudo-inverse*)
 - (b) $\mathbf{A}^+\mathbf{A}\mathbf{A}^+ = \mathbf{A}^+$. (*Outer pseudo-inverse. Any g_1 inverse that satisfies this condition is called a g_2 inverse, or reflexive generalized inverse and is denoted by \mathbf{A}^* .*)
 - (c) $\mathbf{A}^+\mathbf{A}$ is symmetric.
 - (d) $\mathbf{A}\mathbf{A}^+$ is symmetric.
- \mathbf{A}^+ exists and is unique for any matrix \mathbf{A} .
- *Generalized inverse* (or *g_1 inverse*, denoted by \mathbf{A}^- or \mathbf{A}^g): property (a).
- *g_2 inverse* (denoted by \mathbf{A}^*): properties (a)+(b).
- *Moore-Penrose inverse* (denoted by \mathbf{A}^+): properties (a)+(b)+(c)+(d).
- If \mathbf{A} is square and full rank, then the generalized inverse is unique and denoted by \mathbf{A}^{-1} (inverse).
- In practice, the Moore-Penrose inverse \mathbf{A}^+ is easily computed from the singular value decomposition (SVD) of \mathbf{A} .
- $(\mathbf{A}^-)^T$ is a generalized inverse of \mathbf{A}^T .
- $\mathcal{C}(\mathbf{A}) = \mathcal{C}(\mathbf{A}\mathbf{A}^-)$ and $\mathcal{C}(\mathbf{A}^T) = \mathcal{C}((\mathbf{A}^- \mathbf{A})^T)$.
 $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{A}\mathbf{A}^-) = \text{rank}(\mathbf{A}^- \mathbf{A})$.
 “Multiplication by generalized inverse does not change rank.”
- $\text{rank}(\mathbf{A}^-) \geq \text{rank}(\mathbf{A})$. “Generalized inverse has equal or a larger rank than the original matrix.”

System of linear equations

$\mathbf{Ax} = \mathbf{b}$ where $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{b} \in \mathbb{R}^m$.

- When is there a solution? The following statements are equivalent.
 1. The linear system $\mathbf{Ax} = \mathbf{b}$ has a solution (*consistent*)
 2. $\mathbf{b} \in \mathcal{C}(\mathbf{A})$.
 3. $\text{rank}((\mathbf{A}, \mathbf{b})) = \text{rank}(\mathbf{A})$.
 4. $\mathbf{AA}^-\mathbf{b} = \mathbf{b}$.

The last equivalence gives intuition why \mathbf{A}^- is called an inverse.

- What are the solutions to a homogeneous system $\mathbf{Ax} = \mathbf{0}$?
 $\mathcal{N}(\mathbf{A}) = \mathcal{C}(\mathbf{I}_n - \mathbf{A}^-\mathbf{A})$.
- If $\mathbf{Ax} = \mathbf{b}$ is consistent, then $\tilde{\mathbf{x}}$ is a solution to $\mathbf{Ax} = \mathbf{b}$ if and only if

$$\tilde{\mathbf{x}} = \mathbf{A}^-\mathbf{b} + (\mathbf{I}_n - \mathbf{A}^-\mathbf{A})\mathbf{q}$$

for some $\mathbf{q} \in \mathbb{R}^n$.

Interpretation: “a specific solution” + “a vector in the null space of \mathbf{A} ”.

- $\mathbf{Ax} = \mathbf{b}$ is consistent for *all* \mathbf{b} if and only if \mathbf{A} has full row rank.
- If a system is consistent, its solution is unique if and only if \mathbf{A} has full column rank.
- If \mathbf{A} has full row and column rank, then \mathbf{A} is non-singular and the unique solution is $\mathbf{A}^{-1}\mathbf{b}$.

Gramian matrix $\mathbf{A}^T\mathbf{A}$

- $\mathbf{A}^T\mathbf{A}$ is symmetric and positive semidefinite.
- $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{A}^T) = \text{rank}(\mathbf{A}^T\mathbf{A}) = \text{rank}(\mathbf{A}\mathbf{A}^T)$.
- $\mathbf{A}^T\mathbf{A} = \mathbf{0}$ if and only if $\mathbf{A} = \mathbf{0}$.
- $\mathbf{BA}^T\mathbf{A} = \mathbf{CA}^T\mathbf{A}$ if and only if $\mathbf{BA}^T = \mathbf{CA}^T$.

- $\mathbf{A}^\top \mathbf{A} \mathbf{B} = \mathbf{A}^\top \mathbf{A} \mathbf{C}$ if and only if $\mathbf{A} \mathbf{B} = \mathbf{A} \mathbf{C}$.
- For any generalized inverse $(\mathbf{A}^\top \mathbf{A})^-$, $[(\mathbf{A}^\top \mathbf{A})^-]^\top$ is also a generalized inverse of $\mathbf{A}^\top \mathbf{A}$. Note $(\mathbf{A}^\top \mathbf{A})^-$ is not necessarily symmetric.
- $(\mathbf{A}^\top \mathbf{A})^- \mathbf{A}^\top$ is a generalized inverse of \mathbf{A} .
- $\mathbf{A} \mathbf{A}^+ = \mathbf{A} (\mathbf{A}^\top \mathbf{A})^- \mathbf{A}^\top$, where \mathbf{A}^+ is the Moore-Penrose inverse of \mathbf{A} .
- $\mathbf{P}_\mathbf{A} = \mathbf{A} (\mathbf{A}^\top \mathbf{A})^- \mathbf{A}^\top$ is symmetric, idempotent, invariant to the choice of generalized inverse $(\mathbf{A}^\top \mathbf{A})^-$, and projects onto $\mathcal{C}(\mathbf{A})$.

Idempotent matrix and projection

Assume $\mathbf{P} \in \mathbb{R}^{n \times n}$.

- A matrix $\mathbf{P} \in \mathbb{R}^{n \times n}$ is *idempotent* if and only if $\mathbf{P}^2 = \mathbf{P}$.
- A matrix \mathbf{P} is a *projection* on a vector space \mathcal{V} if (a) \mathbf{P} is idempotent, (b) $\mathbf{P} \mathbf{x} \in \mathcal{V}$ for all \mathbf{x} , and (c) $\mathbf{P} \mathbf{z} = \mathbf{z}$ for all $\mathbf{z} \in \mathcal{V}$.
- An idempotent matrix \mathbf{P} is a projection onto $\mathcal{C}(\mathbf{P})$.
- For a *general* matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, the matrices $\mathbf{A} \mathbf{A}^-$ are projections onto $\mathcal{C}(\mathbf{A})$ and $\mathbf{I}_n - \mathbf{A}^- \mathbf{A}$ are projections onto $\mathcal{N}(\mathbf{A})$.




Figure A.1: Three projections onto a column space.

Symmetric idempotent matrix and orthogonal projection

Assume $\mathbf{A} \in \mathbb{R}^{n \times n}$.

- A symmetric, idempotent matrix is called an *orthogonal projection*.
- An orthogonal projection \mathbf{P} satisfies $\mathbf{y} - \mathbf{P}\mathbf{y} \perp \mathbf{v}$ for all $\mathbf{v} \in \mathcal{C}(\mathbf{P})$.
- The orthogonal projection onto a vector space is unique.
- If a symmetric, idempotent matrix \mathbf{P} projects onto \mathcal{V} , then $\mathbf{I} - \mathbf{P}$ projects onto the orthogonal complement \mathcal{V}^\perp .
- Pythagorean theorem: For \mathbf{P} an orthogonal projection,

$$\|\mathbf{y}\|_2^2 = \|\mathbf{P}\mathbf{y}\|_2^2 + \|(\mathbf{I} - \mathbf{P})\mathbf{y}\|_2^2.$$

- Many books use the term “projection” in the sense of orthogonal projection.


Method of least squares

- Goal: Approximate $\mathbf{y} \in \mathbb{R}^n$ by a linear combination of columns of $\mathbf{X} \in \mathbb{R}^{n \times p}$.
- Least squares criterion: $\min Q(\mathbf{b}) = \|\mathbf{y} - \mathbf{X}\mathbf{b}\|_2^2$.
- Any solution to the normal equation $\mathbf{X}^T \mathbf{X} \mathbf{b} = \mathbf{X}^T \mathbf{y}$ (always consistent) is a minimizer of the least squares criterion $Q(\mathbf{b})$.
- Solutions to the normal equation:

$$\hat{\mathbf{b}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y} + (\mathbf{I}_p - (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{X}) \mathbf{q},$$

where $\mathbf{q} \in \mathbb{R}^p$ is arbitrary.

- $(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T$ is a generalized inverse of \mathbf{X} . Therefore the least squares solution applies even when the system $\mathbf{X}\mathbf{b} = \mathbf{y}$ is inconsistent.
- Least squares solution is unique if and only if \mathbf{X} has full column rank.
- $\mathbf{P}_X = \mathbf{X}(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T$ is the orthogonal projection onto $\mathcal{C}(\mathbf{X})$.
- Geometry: The fitted value from the least squares solution $\hat{\mathbf{y}} = \mathbf{P}_X \mathbf{y}$ is the orthogonal projection of the response vector \mathbf{y} onto the column space $\mathcal{C}(\mathbf{X})$.



- $\mathbf{I}_n - \mathbf{P}_X$ is the orthogonal projection onto $\mathcal{N}(\mathbf{X}^T)$.

- Decomposition of \mathbf{y} :

$$\mathbf{y} = \mathbf{P}_X \mathbf{y} + (\mathbf{I}_n - \mathbf{P}_X) \mathbf{y} = \hat{\mathbf{y}} + \hat{\mathbf{e}},$$

where $\hat{\mathbf{y}} \perp \hat{\mathbf{e}}$ and

$$\|\mathbf{y}\|_2^2 = \|\hat{\mathbf{y}}\|_2^2 + \|\hat{\mathbf{e}}\|_2^2.$$

Eigenvalues and eigenvectors

KL chapter 8. Assume $\mathbf{A} \in \mathbb{R}^{n \times n}$ a square matrix.

- *Eigenvalues* are defined as roots of the characteristic equation $\det(\lambda \mathbf{I}_n - \mathbf{A}) = 0$.
- If λ is an eigenvalue of \mathbf{A} , then there exist non-zero $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ such that $\mathbf{A}\mathbf{x} = \lambda\mathbf{x}$ and $\mathbf{y}^T \mathbf{A} = \lambda\mathbf{y}^T$. \mathbf{x} and \mathbf{y} are called the (column) *eigenvector* and *row eigenvector* of \mathbf{A} associated with the eigenvalue λ .
- \mathbf{A} is singular if and only if it has at least one 0 eigenvalue.
- Eigenvectors associated with distinct eigenvalues are linearly independent.
- Eigenvalues of an upper or lower triangular matrix are its diagonal entries: $\lambda_i = a_{ii}$.
- Eigenvalues of an idempotent matrix are either 0 or 1.
- Eigenvalues of an orthogonal matrix have complex modulus 1.
- In most statistical applications, we deal with eigenvalues/eigenvectors of symmetric matrices.
The eigenvalues and eigenvectors of a real *symmetric* matrix are real.
- Eigenvectors associated with distinct eigenvalues of a symmetry matrix are orthogonal.
- Eigen-decomposition of a symmetric matrix: $\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$, where
 - $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_n)$
 - columns of \mathbf{U} are the eigenvectors which are (or can be chosen to be) mutually orthonormal

- A real symmetric matrix is positive semidefinite (positive definite) if and only if all eigenvalues are nonnegative (positive).
- *Spectral radius* $\rho(\mathbf{A}) = \max_i |\lambda_i|$.
- $\mathbf{A} \in \mathbb{R}^{n \times n}$ a square matrix (not required to be symmetric), then $\text{tr}(\mathbf{A}) = \sum_i \lambda_i$ and $|\mathbf{A}| = \prod_i \lambda_i$.

Singular value decomposition

KL chapter 9. Assume $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $p = \min\{m, n\}$.

- Singular value decomposition (SVD): $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$, where
 - $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_m) \in \mathbb{R}^{m \times m}$ is orthogonal
 - $\mathbf{V} = (\mathbf{v}_1, \dots, \mathbf{v}_n) \in \mathbb{R}^{n \times n}$ is orthogonal
 - $\mathbf{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_p) \in \mathbb{R}^{m \times n}$, $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_p \geq 0$.

σ_i are called the *singular values*, \mathbf{u}_i are the *left singular vectors*, and \mathbf{v}_i are the *right singular vectors*.

- $\mathbf{A}\mathbf{v}_i = \sigma_i\mathbf{u}_i$ and $\mathbf{A}^\top\mathbf{u}_i = \sigma_i\mathbf{v}_i$ for $i = 1, \dots, p$.
- Thin SVD. Assume $m \geq n$. \mathbf{A} can be factored as $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$, where
 - $\mathbf{U} \in \mathbb{R}^{m \times n}$, $\mathbf{U}^\top\mathbf{U} = \mathbf{I}_n$
 - $\mathbf{V} \in \mathbb{R}^{n \times n}$, $\mathbf{V}^\top\mathbf{V} = \mathbf{I}_n$
 - $\mathbf{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_n)$

- Relation to eigen-decomposition. Using thin SVD,

$$\begin{aligned}\mathbf{A}^\top\mathbf{A} &= \mathbf{V}\mathbf{\Sigma}\mathbf{U}^\top\mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top = \mathbf{V}\mathbf{\Sigma}^2\mathbf{V}^\top \\ \mathbf{A}\mathbf{A}^\top &= \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top\mathbf{V}\mathbf{\Sigma}\mathbf{U}^\top = \mathbf{U}\mathbf{\Sigma}^2\mathbf{U}^\top.\end{aligned}$$

- Another relation to eigen-decomposition. Using thin SVD,

$$\begin{pmatrix} \mathbf{0}_{n \times n} & \mathbf{A}^\top \\ \mathbf{A} & \mathbf{0}_{m \times m} \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} \mathbf{V} & \mathbf{V} \\ \mathbf{U} & -\mathbf{U} \end{pmatrix} \begin{pmatrix} \mathbf{\Sigma} & \mathbf{0}_{n \times n} \\ \mathbf{0}_{n \times n} & -\mathbf{\Sigma} \end{pmatrix} \frac{1}{\sqrt{2}} \begin{pmatrix} \mathbf{V}^\top & \mathbf{U}^\top \\ \mathbf{V}^\top & -\mathbf{U}^\top \end{pmatrix}.$$

Hence any symmetric eigen-solver can produce the SVD of a matrix \mathbf{A} without forming $\mathbf{A}\mathbf{A}^\top$ or $\mathbf{A}^\top\mathbf{A}$.

- Yet another relation to eigen-decomposition: If the eigendecomposition of a real symmetric matrix is $\mathbf{A} = \mathbf{W}\mathbf{\Lambda}\mathbf{W}^T = \mathbf{W}\text{diag}(\lambda_1, \dots, \lambda_n)\mathbf{W}^T$, then

$$\mathbf{A} = \mathbf{W}\mathbf{\Lambda}\mathbf{W}^T = \mathbf{W} \begin{pmatrix} |\lambda_1| & & \\ & \ddots & \\ & & |\lambda_n| \end{pmatrix} \begin{pmatrix} \text{sgn}(\lambda_1) & & \\ & \ddots & \\ & & \text{sgn}(\lambda_n) \end{pmatrix} \mathbf{W}^T$$

is the SVD of \mathbf{A} .

- Relation to the Moore-Penrose (MP) inverse: Using thin SVD,

$$\mathbf{A}^+ = \mathbf{V}\mathbf{\Sigma}^+\mathbf{U}^T,$$

where $\mathbf{\Sigma}^+ = \text{diag}(\sigma_1^{-1}, \dots, \sigma_r^{-1}, 0, \dots, 0)$, $r = \text{rank}(\mathbf{A})$.

- Denote $\sigma(\mathbf{A}) = (\sigma_1, \dots, \sigma_p)$. Then
 - $\text{rank}(\mathbf{A}) = \# \text{ nonzero singular values} = \|\sigma(\mathbf{A})\|_0$
 - $\mathbf{A} = \mathbf{U}_r\mathbf{\Sigma}_r\mathbf{V}_r^T = \sum_{i=1}^r \sigma_i \mathbf{u}_i \mathbf{v}_i^T$
 - $\|\mathbf{A}\|_F = (\sum_{i=1}^p \sigma_i^2)^{1/2} = \|\sigma(\mathbf{A})\|_2$
 - $\|\mathbf{A}\|_2 = \sigma_1 = \|\sigma(\mathbf{A})\|_\infty$
- Assume $\text{rank}(\mathbf{A}) = r$ and partition $\mathbf{U} = (\mathbf{U}_r, \tilde{\mathbf{U}}_r) \in \mathbb{R}^{m \times m}$ and $\mathbf{V} = (\mathbf{V}_r, \tilde{\mathbf{V}}_r) \in \mathbb{R}^{n \times n}$, then
 - $\mathcal{C}(\mathbf{A}) = \text{span}\{\mathbf{u}_1, \dots, \mathbf{u}_r\}$, $\mathcal{N}(\mathbf{A}^T) = \text{span}\{\mathbf{u}_{r+1}, \dots, \mathbf{u}_m\}$
 - $\mathcal{N}(\mathbf{A}) = \text{span}\{\mathbf{v}_{r+1}, \dots, \mathbf{v}_n\}$, $\mathcal{C}(\mathbf{A}^T) = \text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_r\}$
 - $\mathbf{U}_r\mathbf{U}_r^T$ is the orthogonal projection onto $\mathcal{C}(\mathbf{A})$
 - $\tilde{\mathbf{U}}_r\tilde{\mathbf{U}}_r^T$ is the orthogonal projection onto $\mathcal{N}(\mathbf{A}^T)$
 - $\mathbf{V}_r\mathbf{V}_r^T$ is the orthogonal projection onto $\mathcal{C}(\mathbf{A}^T)$
 - $\tilde{\mathbf{V}}_r\tilde{\mathbf{V}}_r^T$ is the orthogonal projection onto $\mathcal{N}(\mathbf{A})$

Preliminaries of numerical linear algebra

Numerical linear algebra concerns how matrix/vector computations are done in computer. We first look at some basic linear algebra operations.

Flop counts of some basic linear algebra subroutines (BLAS)

See <http://www.netlib.org/blas/> for a complete listing of BLAS functions.

Level	Operation	Name	Dimension	Flops
1	$\alpha \leftarrow \mathbf{x}^T \mathbf{y}$	dot product	$\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$	n
	$\mathbf{y} \leftarrow \mathbf{y} + a\mathbf{x}$	saxpy	$a \in \mathbb{R}, \mathbf{x}, \mathbf{y} \in \mathbb{R}^n$	n
2	$\mathbf{y} \leftarrow \mathbf{y} + \mathbf{A}\mathbf{x}$	gaxpy	$\mathbf{A} \in \mathbb{R}^{m \times n}, \mathbf{x} \in \mathbb{R}^n, \mathbf{y} \in \mathbb{R}^m$	mn
	$\mathbf{A} \leftarrow \mathbf{A} + \mathbf{y}\mathbf{x}^T$	rank one update	$\mathbf{A} \in \mathbb{R}^{m \times n}, \mathbf{x} \in \mathbb{R}^n, \mathbf{y} \in \mathbb{R}^m$	mn
3	$\mathbf{C} \leftarrow \mathbf{C} + \mathbf{A}\mathbf{B}$	matrix multiplication	$\mathbf{A} \in \mathbb{R}^{m \times p}, \mathbf{B} \in \mathbb{R}^{p \times n}, \mathbf{C} \in \mathbb{R}^{m \times n}$	mnp
	$\mathbf{A} \leftarrow \mathbf{A}\mathbf{D}$	column scaling	$\mathbf{A} \in \mathbb{R}^{m \times n}, \mathbf{D} = \text{diag}(d_1, \dots, d_n)$	mn
	$\mathbf{A} \leftarrow \mathbf{D}\mathbf{A}$	row scaling	$\mathbf{A} \in \mathbb{R}^{m \times n}, \mathbf{D} = \text{diag}(d_1, \dots, d_m)$	mn

- Go over the “mxmult” session in R. Also read JM 3.8.
<http://hua-zhou.github.io/teaching/st758-2014fall/mxmult.html>
 Different ways to compute $\mathbf{X}^T \mathbf{W}^{-1} \mathbf{X}$, such as in the weighted least squares.
 $\mathbf{X} \in \mathbb{R}^{n \times p}, \mathbf{W} = \text{diag}(w_1, \dots, w_n) \in \mathbb{R}^{n \times n}$.
 1. `t(X) %*% solve(W) %*% X`: $O(n^3 + n^2p + np^2)$ flops, why need to do the expensive matrix inversion?
 2. `t(X) %*% diag(1 / w) %*% X`: $O(n^2p + np^2)$ flops, why need to save a diagonal matrix and do an extra matrix multiplication?
 3. `(t(X) / w) %*% X`: wrong! \mathbf{w} recycled incorrectly
 4. `t(X) %*% (X / w)`: $O(np^2 + np) = O(np^2)$ flops
 5. `crossprod(X, X / w)`: same as 4, skip the transpose operation
- Another example: Fisher information matrix of a generalized linear model (GLM): $\mathbf{X}^T \mathbf{W} \mathbf{X}$, where $\mathbf{X} \in \mathbb{R}^{n \times p}$ and $\mathbf{W} = \text{diag}(w_1, \dots, w_n)$ are the observation weights.
- Bottom line: Always be flop-aware when writing code.

The form of a mathematical expression and the way the expression should be evaluated in actual practice may be quite different.

- But, for high-performance matrix commutations, it is *not* enough to minimize flops. Pipelining, effective use of memory hierarchy, data layout in memory, ... play important role too.

Vector computer

- Most modern computers are vector machines, which perform vector calculations (saxpy, inner product) fast.
- Vector processing by *pipelining*. E.g., vector addition $z = x + y$




FIG. 1.4.1 *A 3-Cycle Adder*



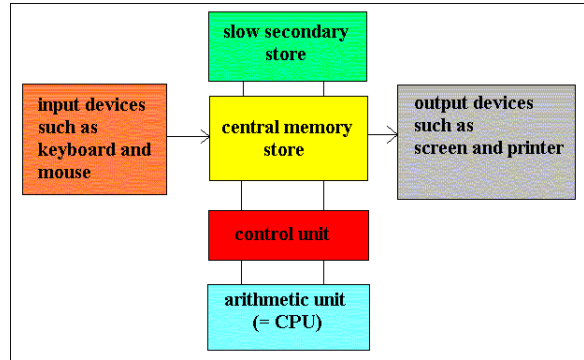



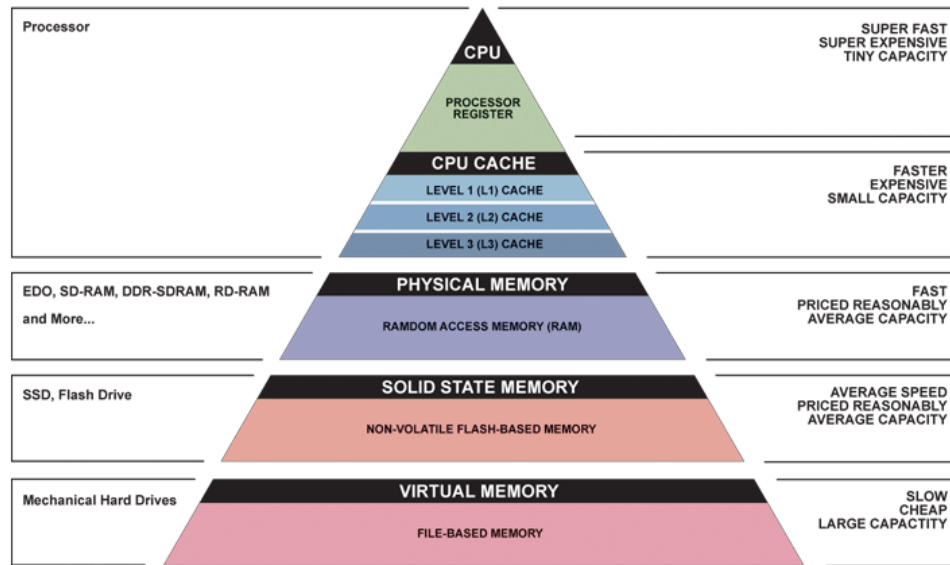
FIG. 1.4.2 *Pipelined Addition*

- One implication of the pipelining technology is that we need to ship vectors to the pipeline fast enough to keep the arithmetic units (ALU) busy and maintain high throughput.

Computer architecture



- Memory hierarchy:



▲ Simplified Computer Memory Hierarchy
Illustration: Ryan J. Leng

Upper the hierarchy, faster the memory accessing speed, and more expensive the memory units.

Key to high performance is effective use of memory hierarchy. True on all architectures.

- Can we keep the super fast arithmetic units busy with enough deliveries of matrix data and ship the results to memory fast enough to avoid backlog?
Answer: use high-level BLAS as much as possible
- Why high-level BLAS?

BLAS	Dimension	Mem Refs	Flops	Ratio
Level 1: $\mathbf{y} \leftarrow \mathbf{y} + \mathbf{ax}$	$\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$	$3n$	n	3:1
Level 2: $\mathbf{y} \leftarrow \mathbf{y} + \mathbf{Ax}$	$\mathbf{x}, \mathbf{y} \in \mathbb{R}^n, \mathbf{A} \in \mathbb{R}^{n \times n}$	n^2	n^2	1:1
Level 3: $\mathbf{C} \leftarrow \mathbf{C} + \mathbf{AB}$	$\mathbf{A}, \mathbf{B}, \mathbf{C} \in \mathbb{R}^{n \times n}$	$4n^2$	n^3	4:n

- BLAS 1 tend to be *memory bandwidth-limited*. E.g., Xeon X5650 CPU has a theoretical throughput of 128 DP GFLOPS but a max memory bandwidth of 32GB/s.
- Higher level BLAS (3 or 2) make more effective use of ALUs (keep them busy).

- Message: Although we state many algorithms (solving linear equations, least squares, eigen-decomposition, SVD, ...) in terms of inner product and saxpy, the actual implementation may be quite different.
- A distinction between LAPACK and LINPACK is that LAPACK makes use of higher level BLAS as much as possible (usually by smart partitioning) to increase the so-called *level-3 fraction*.

5 Lecture 5, Sep 11

Announcements

- TA office hours: Fri @ 10A-12P
- HW1 due today. Submit hardcopy and email R or Rmd code
- HW2 posted, due Tue Sep 23

Last time

- Flop counts of some BLAS subroutines
- Memory hierarchy and its implication for implementation of BLAS

Today

- Numerical linear algebra preliminaries: effects of data layout
- Numerical linear algebra: GE, LU

Effect of data layout

- Data layout in memory effects execution speed too. It is much faster to move chunks of data in memory than retrieving/writing scattered data.
- Storage mode: column-major (FORTRAN, MATLAB, R) vs row-major (C/C++).
- Take matrix multiplication as an example. Assume the storage is column-major, such as in FORTRAN. $\mathbf{C} \leftarrow \mathbf{C} + \mathbf{A}\mathbf{B}$, where $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{B} \in \mathbb{R}^{n \times p}$, $\mathbf{C} \in \mathbb{R}^{m \times p}$. There are 6 variants of the algorithms according to the order in the triple loops. We pay attention to the innermost loop, where the vector calculation occurs,

```

jki or kji:      for  $i = 1:m$ 
                        $C(i,j) = C(i,j) + A(i,k)B(k,j)$ 
                       end

ikj or kij:      for  $j = 1:n$ 
                        $C(i,j) = C(i,j) + A(i,k)B(k,j)$ 
                       end

ijk or jik:      for  $k = 1:p$ 
                        $C(i,j) = C(i,j) + A(i,k)B(k,j)$ 
                       end

```

and the associated *stride* when accessing the three matrices in memory (assuming column-major storage)

Variant	A Stride	B Stride	C Stride
<i>jki</i> or <i>kji</i>	Unit	0	Unit
<i>ikj</i> or <i>kij</i>	0	Non-Unit	Non-Unit
<i>ijk</i> or <i>jik</i>	Non-Unit	Unit	0

Apparently the variants *jki* or *kji* are preferred.

- Message: data storage mode effects algorithm implementation too.

Solving linear equations

We consider algorithms for solving linear equations $\mathbf{Ax} = \mathbf{b}$, a ubiquitous task in statistics. Idea: turning original problem into an “easy” one, e.g., triangular system.

Triangular system

- *Forward substitution* to solve $\mathbf{Lx} = \mathbf{b}$, where $\mathbf{L} \in \mathbb{R}^{n \times n}$ is lower triangular

$$\begin{bmatrix} a_{11} & 0 & \dots & 0 \\ a_{21} & a_{22} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & \dots & a_{mm} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_m \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_m \end{bmatrix}$$

$$x_1 = b_1/a_{11}$$

$$x_2 = (b_2 - a_{21}x_1)/a_{22}$$

$$x_3 = (b_3 - a_{31}x_1 - a_{32}x_2)/a_{33}$$

⋮

$$x_m = (b_m - a_{m1}x_1 - a_{m2}x_2 - \dots - a_{m,m-1}x_{m-1})/a_{mm}$$

$n^2/2$ flops ($n^2/2$ multiplications/divisions and $n^2/2$ additions/subtractions)
and \mathbf{L} is accessed by row.

- *Back substitution* to solve $\mathbf{U}\mathbf{x} = \mathbf{b}$ where $\mathbf{U} \in \mathbb{R}^{n \times n}$ is upper triangular

$$\begin{bmatrix} a_{11} & \dots & a_{1,m-1} & a_{1m} \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \dots & a_{m-1,m-1} & a_{m-1,m} \\ 0 & \dots & 0 & a_{mm} \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_{m-1} \\ x_m \end{bmatrix} = \begin{bmatrix} b_1 \\ \vdots \\ b_{m-1} \\ b_m \end{bmatrix}$$

$$x_m = b_m/a_{mm}$$

$$x_{m-1} = (b_{m-1} - a_{m-1,m}x_m)/a_{m-1,m-1}$$

$$x_{m-2} = (b_{m-2} - a_{m-2,m-1}x_{m-1} - a_{m-2,m}x_m)/a_{m-2,m-2}$$

⋮

$$x_1 = (b_1 - a_{12}x_2 - a_{13}x_3 - \dots - a_{1m}x_m)/a_{11}$$

$n^2/2$ flops ($n^2/2$ multiplications/divisions and $n^2/2$ additions/subtractions)
and \mathbf{U} is accessed by row.

- Column version: reverse the order of looping.
- BLAS level 2 function: `?trsv` (triangular solve with one right hand side)
- BLAS level 3 function: `?trsm` (matrix triangular solve, i.e., multiple right hand sides)

- In R, `forwardsolve()` and `backsolve()` (wrappers of `dtrsm`)
- Eigenvalues of \mathbf{L} are diagonal entries $\lambda_i = \ell_{ii}$. $\det(\mathbf{L}) = \prod_i \ell_{ii}$.
- A *unit triangular matrix* is a triangular matrix with all diagonal entries being 1.
- The algebra of triangular matrices (HW2)
 - The product of two upper (lower) triangular matrices is upper (lower) triangular.
 - The inverse of an upper (lower) triangular matrix is upper (lower) triangular.
 - The product of two unit upper (lower) triangular matrices is unit upper (lower) triangular.
 - The inverse of a unit upper (lower) triangular matrix is unit upper (lower) triangular.

Gaussian elimination and LU decomposition

Given a system of linear algebraic equations

$$\begin{bmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ a_{21} & a_{22} & \dots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & \dots & a_{mn} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_m \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_m \end{bmatrix}$$

Step 1: Each row times a_{11}/a_{k1} ,
then use row one to subtract other rows.

$$\Rightarrow \begin{bmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ 0 & \tilde{a}_{22} & \dots & \tilde{a}_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \tilde{a}_{n2} & \dots & \tilde{a}_{nn} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} b_1 \\ \tilde{b}_2 \\ \vdots \\ \tilde{b}_n \end{bmatrix}$$

Step 2: The second row and down multiply by $\tilde{a}_{22}/\tilde{a}_{k2}$,
then use row two to subtract every row below.

$$\Rightarrow \begin{bmatrix} a_{11} & a_{12} & a_{13} & \dots & a_{1n} \\ 0 & \tilde{a}_{22} & \tilde{a}_{23} & \dots & \tilde{a}_{2n} \\ 0 & 0 & \tilde{a}_{33} & \dots & \tilde{a}_{3n} \\ \vdots & \vdots & \dots & \ddots & \vdots \\ 0 & 0 & \tilde{a}_{n3} & \dots & \tilde{a}_{nn} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} b_1 \\ \tilde{b}_2 \\ \tilde{b}_3 \\ \vdots \\ \tilde{b}_n \end{bmatrix}$$

Step 3: Similar to the previous two steps, repeat until all
elements in the lower triangle of the matrix A
become zeros.

$$\Rightarrow \begin{bmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ 0 & \tilde{a}_{22} & \dots & \tilde{a}_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \vdots & \vdots & \dots & \tilde{a}_{nn} \\ 0 & 0 & \dots & \tilde{a}_{nn} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} b_1 \\ \tilde{b}_2 \\ \vdots \\ \tilde{b}_n \end{bmatrix}$$

- Solve $\mathbf{Ax} = \mathbf{b}$ for a general matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$.
- An important contribution by Gauss. No linear algebra in 1800!
- Idea: a series of *elementary operations* that turn \mathbf{A} into a triangular system.
- We consider the square \mathbf{A} case first.

- *Elementary operator matrix* $\mathbf{E}_{jk}(c)$ is the identity with the 0 in position (j, k) replaced by c . For any vector \mathbf{x} ,

$$\mathbf{E}_{jk}(c)\mathbf{x} = (x_1, \dots, x_{j-1}, x_j + cx_k, x_{j+1}, \dots, x_n)^\top.$$

Applying $\mathbf{E}_{jk}(c)$ to both sides of the system $\mathbf{Ax} = \mathbf{b}$ replaces the j -th equation $\mathbf{a}_{j*}^\top \mathbf{x} = b_j$ by $\mathbf{a}_{j*}^\top \mathbf{x} + c\mathbf{a}_{k*}^\top \mathbf{x} = b_j + cb_k$. For $j > k$, $\mathbf{E}_{jk}(c) = \mathbf{I} + c\mathbf{e}_j\mathbf{e}_k^\top$ is unit lower triangular and full rank. $\mathbf{E}_{jk}^{-1}(c) = \mathbf{E}_{jk}(-c)$.

- Zeroing the first column

$$\begin{aligned} \mathbf{E}_{21}(c_2^{(1)})\mathbf{Ax} &= \mathbf{E}_{21}(c_2^{(1)})\mathbf{b} \\ \mathbf{E}_{31}(c_3^{(1)})\mathbf{E}_{21}(c_2^{(1)})\mathbf{Ax} &= \mathbf{E}_{31}(c_3^{(1)})\mathbf{E}_{21}(c_2^{(1)})\mathbf{b} \\ &\vdots \\ \mathbf{E}_{n1}(c_n^{(1)})\cdots\mathbf{E}_{31}(c_3^{(1)})\mathbf{E}_{21}(c_2^{(1)})\mathbf{Ax} &= \mathbf{E}_{n1}(c_n^{(1)})\cdots\mathbf{E}_{31}(c_3^{(1)})\mathbf{E}_{21}(c_2^{(1)})\mathbf{b} \end{aligned}$$

where $c_i^{(1)} = -a_{i1}/a_{11}$. Denote $\mathbf{M}_1 = \mathbf{E}_{n1}(c_n^{(1)})\cdots\mathbf{E}_{31}(c_3^{(1)})\mathbf{E}_{21}(c_2^{(1)})$.

- Then zero the k -th column for $k = 2, \dots, n-1$ sequentially. This results in a transformed linear system $\mathbf{Ux} = \tilde{\mathbf{b}}$, where $\mathbf{U} = \mathbf{M}_{n-1}\cdots\mathbf{M}_1\mathbf{A}$ is upper triangular and $\tilde{\mathbf{b}} = \mathbf{M}_{n-1}\cdots\mathbf{M}_1\mathbf{b}$. \mathbf{M}_k has the shape

$$\mathbf{M}_k = \mathbf{E}_{n,k}^{(k)} \cdots \mathbf{E}_{k+1,k}^{(k)} = \begin{pmatrix} 1 & & & & & \\ & \ddots & & & & \\ & & 1 & & & \\ & & c_{k+1}^{(k)} & 1 & & \\ & & \vdots & & \ddots & \\ & & c_n^{(k)} & & & 1 \end{pmatrix},$$

where $c_i^{(k)} = -\tilde{a}_{ik}^{(k-1)}/\tilde{a}_{kk}^{(k-1)}$. \mathbf{M}_k is unit lower triangular and full rank. \mathbf{M}_k are called the *Gauss transformations*.

- Let $\mathbf{L} = \mathbf{M}_1^{-1}\cdots\mathbf{M}_{n-1}^{-1}$. We have the decomposition

$$\mathbf{A} = \mathbf{LU}.$$

\mathbf{M}_k is unit upper triangular, so \mathbf{M}_k^{-1} and thus \mathbf{L} is unit lower triangular.

- Where is \mathbf{L} ? Note $\mathbf{M}_k = \mathbf{I} + (0, \dots, 0, c_{k+1}^{(k)}, \dots, c_n^{(k)})^\top \mathbf{e}_k^\top$. By Sherman-Morrison, $\mathbf{M}_k^{-1} = \mathbf{I} - (0, \dots, 0, c_{k+1}^{(k)}, \dots, c_n^{(k)})^\top \mathbf{e}_k^\top$. So the entries of \mathbf{L} are simply $\ell_{ik} = -c_i^{(k)}$, $i > k$, the negative of multipliers in GE.
- The whole LU procedure is done in place, i.e., \mathbf{A} is overwritten by \mathbf{L} and \mathbf{U} .
- Implementation: outer product LU (kij loop), block outer product LU (higher level-3 fraction), Crout's algorithm (jki loop), ...

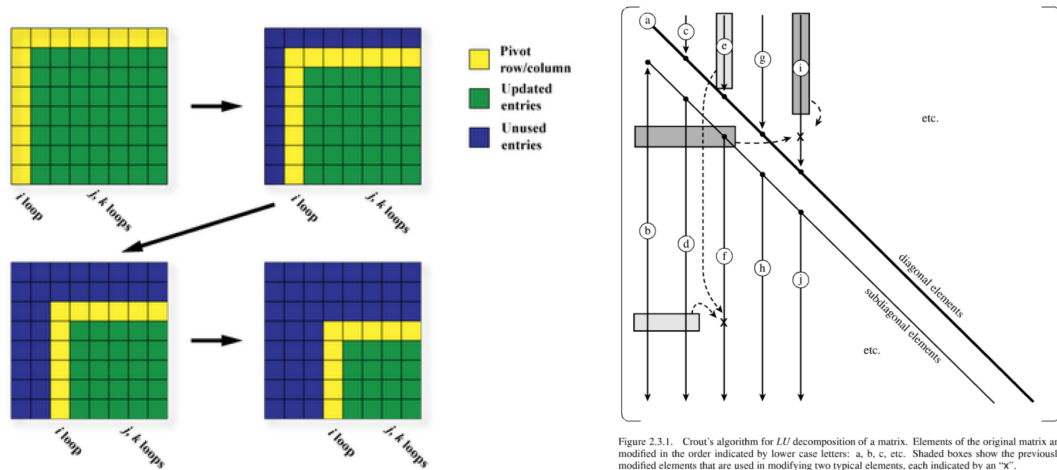


Figure 2.3.1. Crout's algorithm for LU decomposition of a matrix. Elements of the original matrix are modified in the order indicated by lower case letters: a, b, c, etc. Shaded boxes show the previously modified elements that are used in modifying two typical elements, each indicated by an "X".

- LU decomposition exists if the principal sub-matrix $A(1 : k, 1 : k)$ is non-singular for $k = 1, \dots, n - 1$. If the LU decomposition exists and \mathbf{A} is non-singular, then the LU decomposition is unique and $\det(\mathbf{A}) = \prod_{i=1}^n u_{ii}$.
- This *forward elimination* or *LU decomposition* costs $(n-1)^2 + (n-2)^2 + \dots + 1^2 \approx \frac{1}{3}n^3$ flops ($n^3/3$ multiplications and $n^3/3$ additions).
- Given LU, one right hand side costs n^2 flops (one backward substitution and then forward substitution).
- For inversion, there are n right hand sides \mathbf{e}_i . However, taking advantage of zeros reduces n^3 flops to $\frac{2}{3}n^3$ (see JM exercise 3.2). So matrix inversion costs $\frac{1}{3}n^3 + \frac{2}{3}n^3 = n^3$ flops in total.
- We do *not* compute matrix inverse unless (i) it is absolutely necessary to compute s.e., (2) number of right hand sides is much larger than n , (3) n is small.

- LU decomposition of a rectangular matrix $\mathbf{A} \in \mathbf{R}^{m \times n}$ exists if $\mathbf{A}(1 : k, 1 : k)$ is non-singular for $k = \min\{m, n\}$. For example,

$$\begin{pmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 3 & 1 \\ 5 & 2 \end{pmatrix} \begin{pmatrix} 1 & 2 \\ 0 & -2 \end{pmatrix}$$

and

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 4 & 1 \end{pmatrix} \begin{pmatrix} 1 & 2 & 3 \\ 0 & -3 & -6 \end{pmatrix}.$$

Slight modification to the algorithm.

6 Lecture 6, Sep 16

Announcements

- HW1 returned:
 - Change of code naming convention for later homework: `LastFirstHW2.R` or `LastFirstHW2.Rmd`, e.g., `ZhouHuaHW2.Rmd`
 - Identify yourself in your code
 - Submit code on time: late penalty 5pts/day (?)
 - Style: indenting, no more 80 characters per line, space around binary operators, space after comma, ...
 - Q1: work out the computer arithmetic on at least one case, so you're clear about what's happening with rounding, losing significant digits when matching exponents, normalization, ...
 - Q4: `crossprod()`, `tcrossprod()`, `outer()`, `row()`, `col()`, `rowSums()`, `colSums()` functions in R, subsetting in R, ...
 - Q5: vectorize code, determinant from Cholesky decomposition
 - Sketch of solution: <http://hua-zhou.github.io/teaching/st758-2014fall/hw01sol.html>

Last time

- Effect of data layout (column-major vs row-major)
- Triangular system: $n^2/2$ flops for forward substitution or backward substitution
- GE and LU decomposition (JM 3.4)

Today

- GE and LU: pivoting
- Cholesky decomposition (KL 7.7 and JM 3.5)

Pivoting for LU

- What if we encounter a *pivot* $\tilde{a}_{kk}^{(k-1)}$ being 0 or close to 0 due to underflow?
- Think about $\mathbf{A} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. Does it have a solution for arbitrary \mathbf{b} ? Does GE work?
- Work on the example

$$\begin{aligned} 0.0001x_1 + x_2 &= 1 \\ x_1 + x_2 &= 2, \end{aligned}$$

which has solution $x_1 = 1.0001$ and $x_2 = 0.9999$. Suppose we have 3 digits of precision. After first step of elimination, we have (catastrophic cancellation happens)

$$\begin{aligned} 0.0001x_1 + x_2 &= 1 \\ -10,000x_2 &= -10,000 \end{aligned}$$

and the solution by back substitution is $x_2 = 1.000$ and $x_1 = 0.000$.

- Message: zero or very small pivots cause trouble.
Solution: *pivoting*.
- *Partial pivoting*: at the k -th stage the equation with $\max_{i=k}^n |\tilde{a}_{ik}^{(k-1)}|$ is moved into the k -th row. Thus we have $\mathbf{M}_{n-1}\mathbf{P}_{n-1} \cdots \mathbf{M}_1\mathbf{P}_1\mathbf{A} = \mathbf{U}$.
- With partial pivoting, it can be shown that

$$\mathbf{PA} = \mathbf{LU},$$

where $\mathbf{P} = \mathbf{P}_{n-1} \cdots \mathbf{P}_1$, \mathbf{L} is unit lower triangular with $|\ell_{ij}| \leq 1$, and \mathbf{U} is upper triangular.

- $\det(\mathbf{P}) \det(\mathbf{A}) = \det(\mathbf{U}) = \prod_{i=1}^n u_{ii}$.
- To solve $\mathbf{Ax} = \mathbf{b}$, we solve two triangular systems

$$\mathbf{Ly} = \mathbf{Pb} \quad \text{and} \quad \mathbf{Ux} = \mathbf{y},$$

costing n^2 flops.

- *Complete pivoting*: Do both row and column interchanges so that the largest entry in the sub matrix $\mathbf{A}(k : n, k : n)$ is permuted to the (k, k) -th entry. This yields the decomposition $\mathbf{PAQ} = \mathbf{LU}$, where $|\ell_{ij}| \leq 1$.
- Warning: In actual implementation, we do not really need to interchange rows/columns for pivoting. Just keep track of the indices we have interchanged.
- Gaussian elimination with partial pivoting is one of the most commonly used methods for solving general linear systems. Complete pivoting is stable but costs more computation. Partial pivoting is stable most of times.
- LAPACK: ?GETRF does $\mathbf{PA} = \mathbf{LU}$ (LU decomposition with partial pivoting).
- In R, `solve()` implicitly performs LU decomposition (wrapper of LAPACK routine DGESV). `solve()` allows specifying a single or multiple right hand sides. If none, it computes the matrix inverse. The `matrix` package contains `lu()` function.

Cholesky decomposition (symmetric LU)

JM 3.5 and KL 7.7



Archives de l'Ecole polytechnique (Fonds A. Cholesky)

- A basic tenet in numerical analysis:

The structure should be exploited whenever solving a problem.

Common structures include: symmetry, definiteness, sparsity, Kronecker product, low rank, ...

- Consider solving the normal equation $\mathbf{X}^T \mathbf{X} \boldsymbol{\beta} = \mathbf{X}^T \mathbf{y}$ for linear regression. The coefficient matrix $\mathbf{X}^T \mathbf{X}$ is symmetric and positive semidefinite, how to exploit this structure?
- Theorem: Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric and positive definite. Then $\mathbf{A} = \mathbf{L} \mathbf{L}^T$ where \mathbf{L} is lower triangular with positive diagonal entries and is unique.

Proof by induction. If $n = 1$, then $\ell = \sqrt{a}$. For $n > 1$, the block equation

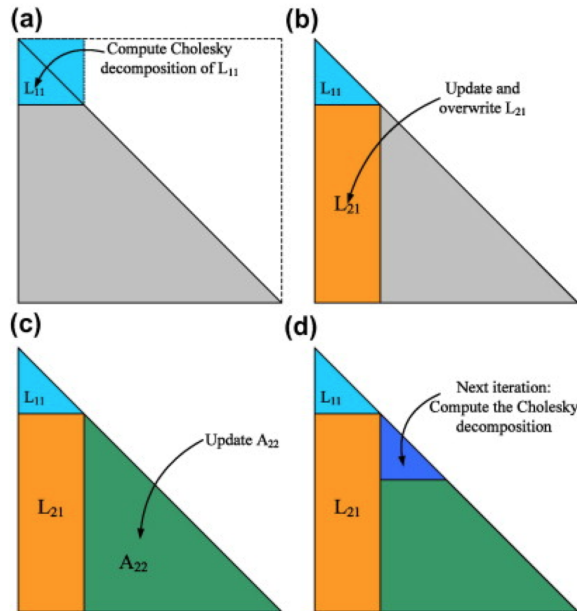
$$\begin{pmatrix} a_{11} & \mathbf{a}^T \\ \mathbf{a} & \mathbf{A}_{22} \end{pmatrix} = \begin{pmatrix} \ell_{11} & \mathbf{0}_{n-1}^T \\ \mathbf{l} & \mathbf{L}_{22} \end{pmatrix} \begin{pmatrix} \ell_{11} & \mathbf{l}^T \\ \mathbf{0}_{n-1} & \mathbf{L}_{22}^T \end{pmatrix}.$$

has solution

$$\begin{aligned} \ell_{11} &= \sqrt{a_{11}} \\ \mathbf{l} &= \ell_{11}^{-1} \mathbf{a} \\ \mathbf{L}_{22} \mathbf{L}_{22}^T &= \mathbf{A}_{22} - \mathbf{l} \mathbf{l}^T = \mathbf{A}_{22} - a_{11}^{-1} \mathbf{a} \mathbf{a}^T. \end{aligned}$$

Now $a_{11} > 0$ (why?), so ℓ_{11} and \mathbf{l} are uniquely determined. $\mathbf{A}_{22} - a_{11}^{-1} \mathbf{a} \mathbf{a}^T$ is positive definite because \mathbf{A} is positive definite (why?). By induction hypothesis, \mathbf{L}_{22} exists and is unique. \square

- The constructive proof completely specifies the algorithm.



- Computational cost: $\frac{1}{2}[(n-1)^2 + (n-2)^2 + \dots + 1^2] \approx \frac{1}{6}n^3$ (half the cost of LU decomposition due to symmetry) plus n square roots.
- Avoid square roots: \mathbf{LDL}^\top decomposition.
- Pivoting? In general Cholesky decomposition is very stable. Failure of the decomposition simply means \mathbf{A} is not positive definite. It is an efficient way to test positive definiteness.

If zero pivots $\tilde{a}_{ii} = 0$ are encountered, we can still continue the algorithm by setting $\ell_{ii} = 0$ and $\mathbf{l} = \mathbf{0}$. A better alternative is to use Cholesky decomposition with symmetric pivoting.

- Any matrix $\mathbf{X} \in \mathbb{R}^{n \times n}$ is a *square root* of $\mathbf{A} \succeq \mathbf{0}_{n \times n}$ if $\mathbf{A} = \mathbf{X}\mathbf{X}$. Note that Cholesky factor is *not* a square root of \mathbf{A} unless \mathbf{A} is diagonal.

7 Lecture 7, Sep 18

Announcement

- HW2 hints: <http://hua-zhou.github.io/teaching/st758-2014fall/st758fall2014/2014/09/16/hw2-hints.html>
- HW3 will be posted today. Due Tue Sep 30.

Last time

- LU with partial pivoting: $\mathbf{PA} = \mathbf{LU}$ (still $n^3/3$ flops)
- Cholesky decomposition $\mathbf{A} = \mathbf{LL}^T$, \mathbf{A} p.d.: $n^3/6$ flops taking advantage of symmetry

Today

- Cholesky decomposition for \mathbf{A} p.s.d.: symmetric pivoting
- Linear regression by Cholesky (JM 3.5, KL 7.7)
- QR decomposition (JM 5.4-5.8, KL 7.8-7.9)

Cholesky with symmetric pivoting

$\mathbf{A} \succeq \mathbf{0}_{n \times n}$ (p.s.d.)

- When \mathbf{A} does not have full rank, e.g., $\mathbf{X}^T \mathbf{X}$ with a non-full column rank \mathbf{X} , we encounter $\tilde{a}_{kk} = 0$ during the procedure.
- *Symmetric pivoting.* At each stage k , we permute both row and column such that $\max_{i=k}^n \tilde{a}_{ik}$ becomes the pivot. If we encounter $\max_{i=k}^n \tilde{a}_{ik} = 0$, then $A(k:n, k:n) = \mathbf{0}$ (why?) and the algorithm terminates.
- With symmetric pivoting: $\mathbf{PAP}^T = \mathbf{LL}^T$, where \mathbf{P} is a permutation matrix and $\mathbf{L} \in \mathbb{R}^{n \times r}$, $r = \text{rank}(\mathbf{A})$.
- In R, `chol()` is a wrapper function for LAPACK routines DPOTRF (p.d.) and DPSTRF (p.s.d. with pivoting).

- Option `pivot = FALSE` calls DPOTRF. It does $\mathbf{A} = \mathbf{R}^T \mathbf{R}$ and gives error message if \mathbf{A} is not full rank.
- Option `pivot = TRUE` calls DPSTRF. It does symmetric pivoting $\mathbf{PAP}^T = \mathbf{R}^T \mathbf{R}$ and yields rank and pivot.
- Option `tol` passes the tolerance to LAPACK for deciding zero pivots. Default is $n \cdot \text{machine epsilon} \cdot \max(\text{diag}(\mathbf{A}))$.

Applications of Cholesky decomposition

There are numerous applications of Cholesky decomposition.

- *No inversion* mentality: Whenever we see matrix inverse, we should think in terms of solving linear equations. If the matrix is positive semidefinite, Cholesky decomposition applies.
- Example: multivariate normal density $N_n(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, $\boldsymbol{\Sigma}$ is p.d.

$$-\frac{n}{2} \ln(2\pi) - \frac{1}{2} \ln \det \boldsymbol{\Sigma} - \frac{1}{2} (\mathbf{y} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{y} - \boldsymbol{\mu}).$$

- Method 1: (a) compute explicit inverse $\boldsymbol{\Sigma}^{-1}$ (n^3 flops), (b) compute quadratic form ($n^2 + n$ flops), (c) compute determinant ($n^3/3$ flops).
- Method 2: (a) Cholesky decomposition $\boldsymbol{\Sigma} = \mathbf{L}\mathbf{L}^T$ ($n^3/6$ flops), (b) Solve $\mathbf{L}\mathbf{x} = \mathbf{y} - \boldsymbol{\mu}$ by forward substitutions ($n^2/2$ flops), (c) compute quadratic form $\mathbf{x}^T \mathbf{x}$ (n flops), and (d) compute determinant from Cholesky factor (n flops).

Which method is better?

- Compute Moore-Penrose inverse \mathbf{A}^+ . (HW3)
- Linear regression.

Linear regression by Cholesky (method of normal equations)

Assume $\mathbf{X} \in \mathbb{R}^{n \times p}$ has full column rank. (For rank deficient \mathbf{X} , use Cholesky with symmetric pivoting.)

- It is easier to work on the augmented matrix

$$\begin{pmatrix} \mathbf{X}^\top \mathbf{X} & \mathbf{X}^\top \mathbf{y} \\ \mathbf{y}^\top \mathbf{X} & \mathbf{y}^\top \mathbf{y} \end{pmatrix} = \begin{pmatrix} \mathbf{L} & \mathbf{0} \\ \mathbf{l}^\top & d \end{pmatrix} \begin{pmatrix} \mathbf{L}^\top & \mathbf{l} \\ \mathbf{0}^\top & d \end{pmatrix} = \begin{pmatrix} \mathbf{L}\mathbf{L}^\top & \mathbf{L}\mathbf{l} \\ \mathbf{l}^\top \mathbf{L}^\top & \|\mathbf{l}\|_2^2 + d^2 \end{pmatrix}.$$

Normal equation implies the equation

$$\mathbf{X}^\top \mathbf{X} \boldsymbol{\beta} = \mathbf{L}\mathbf{L}^\top \boldsymbol{\beta} = \mathbf{X}^\top \mathbf{y} = \mathbf{L}\mathbf{l} \text{ or } \mathbf{L}^\top \boldsymbol{\beta} = \mathbf{l},$$

which we can solve for $\boldsymbol{\beta}$ in $p^2/2$ flops. Since $\mathbf{l} = \mathbf{L}^{-1} \mathbf{X}^\top \mathbf{y}$, we have

$$\mathbf{l}^\top \mathbf{l} = \mathbf{y}^\top \mathbf{X} (\mathbf{L}\mathbf{L}^\top)^{-1} \mathbf{X}^\top \mathbf{y} = \mathbf{y}^\top \mathbf{X} (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y} = \mathbf{y}^\top \mathbf{P}_\mathbf{X} \mathbf{y} = \|\hat{\mathbf{y}}\|_2^2$$

and

$$d^2 = \mathbf{y}^\top \mathbf{y} - \mathbf{l}^\top \mathbf{l} = \mathbf{y}^\top (\mathbf{I} - \mathbf{P}_\mathbf{X}) \mathbf{y} = \|\mathbf{y} - \hat{\mathbf{y}}\|_2^2 = \text{SSE}.$$

If standard errors are needed, we do inversion $(\mathbf{X}^\top \mathbf{X})^{-1} = (\mathbf{L}\mathbf{L}^\top)^{-1} = \mathbf{L}^{-\top} \mathbf{L}^{-1}$.

Use `chol2inv()` in R function for this purpose.

- In summary, linear regression by Cholesky, aka the method of normal equations:

- Form the lower triangular part of $(\mathbf{X}, \mathbf{y})^\top (\mathbf{X}, \mathbf{y})$ ($n(p+1)^2/2$ flops)
- Cholesky decomposition of the augmented system $\begin{pmatrix} \mathbf{X}^\top \mathbf{X} & \mathbf{X}^\top \mathbf{y} \\ \mathbf{y}^\top \mathbf{X} & \mathbf{y}^\top \mathbf{y} \end{pmatrix}$ ($(p+1)^3/6$ flops)
- Solve $\mathbf{L}^\top \boldsymbol{\beta} = \mathbf{l}$ for regression coefficients $\hat{\boldsymbol{\beta}}$ ($p^2/2$ flops)
- If want standard errors, estimate σ^2 by $\hat{\sigma}^2 = d^2/(n-p)$ and compute $\hat{\sigma}^2 (\mathbf{X}^\top \mathbf{X})^{-1} = \hat{\sigma}^2 (\mathbf{L}\mathbf{L}^\top)^{-1}$ ($2p^3/3$ flops)

Total cost is $p^3/6 + np^2/2$ flops (without s.e.) or $5p^3/6 + np^2/2$ flops (with s.e.).

QR decomposition and linear regression

Assume $\mathbf{X} \in \mathbb{R}^{n \times p}$ has full column rank.

- QR decomposition: $\mathbf{X} = \mathbf{Q}\mathbf{R}$, where $\mathbf{Q} \in \mathbb{R}^{n \times n}$, $\mathbf{Q}^\top \mathbf{Q} = \mathbf{I}_n$, and $\mathbf{R} \in \mathbb{R}^{n \times p}$.
 - first p columns of \mathbf{Q} form an orthonormal basis of $\mathcal{C}(\mathbf{X})$
 - last $n-p$ columns of \mathbf{Q} form an orthonormal basis of $\mathcal{N}(\mathbf{X}^\top)$

- (Thin/Skinny QR) Then $\mathbf{X} = \mathbf{Q}_1 \mathbf{R}_1$ where $\mathbf{Q}_1 \in \mathbb{R}^{n \times p}$ has orthogonal columns, $\mathbf{Q}_1^\top \mathbf{Q}_1 = \mathbf{I}_p$, and $\mathbf{R}_1 \in \mathbb{R}^{p \times p}$ is an invertible upper triangular matrix with positive diagonal entries.
- For linear regression, we only need skinny QR.
Note $\mathbf{X}^\top \mathbf{X} = \mathbf{R}_1^\top \mathbf{R}_1$ yields the Cholesky decomposition of $\mathbf{X}^\top \mathbf{X}$.
- Better to perform (skinny) QR on the augmented matrix

$$\begin{pmatrix} \mathbf{X} & \mathbf{y} \end{pmatrix} = \begin{pmatrix} \mathbf{Q} & \mathbf{q} \end{pmatrix} \begin{pmatrix} \mathbf{R} & \mathbf{r} \\ \mathbf{0}_p^\top & d \end{pmatrix} = \begin{pmatrix} \mathbf{QR} & \mathbf{Qr} + d\mathbf{q} \end{pmatrix}.$$

Normal equation $\mathbf{X}^\top \mathbf{X} \boldsymbol{\beta} = \mathbf{X}^\top \mathbf{y}$ implies

$$\mathbf{R} \boldsymbol{\beta} = \mathbf{R}^{-\top} \mathbf{X}^\top \mathbf{y} = \mathbf{R}^{-\top} \mathbf{R}^\top \mathbf{Q}^\top \mathbf{y} = \mathbf{Q}^\top \mathbf{y} = \mathbf{r},$$

which is easy to solve for $\boldsymbol{\beta}$. The fitted value is $\hat{\mathbf{y}} = \mathbf{X} \hat{\boldsymbol{\beta}} = \mathbf{QR} \mathbf{R}^{-1} \mathbf{r} = \mathbf{Qr}$.

The residual is

$$\hat{\mathbf{e}} = \mathbf{y} - \mathbf{X} \hat{\boldsymbol{\beta}} = \mathbf{y} - \mathbf{Qr} = d\mathbf{q}$$

and $\text{SSE} = \|\hat{\mathbf{e}}\|_2^2 = d^2$. The projection matrix is


$$\mathbf{X}(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X} = \mathbf{QR}(\mathbf{R}^\top \mathbf{R})^{-1} \mathbf{R}^\top \mathbf{Q}^\top = \mathbf{QQ}^\top.$$

- Three numerical methods to compute QR: (modified) Gram-Schmidt, Householder transform, (fast) Givens transform


QR by (modified) Gram-Schmidt



Jørgen Pedersen Gram in an undated photo



Erhard Schmidt (courtesy MFO)



Assume $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_p) \in \mathbb{R}^p$ has full column rank.

- Gram-Schmidt algorithm produces the skinny QR.
- Gram-Schmidt algorithm orthonormalizes a set of non-zero, *linearly independent* vectors $\mathbf{x}_1, \dots, \mathbf{x}_p$. Initialize $\mathbf{q}_1 = \mathbf{x}_1 / \|\mathbf{x}_1\|_2$; then for $k = 2, \dots, p$,

$$\mathbf{v}_k = \mathbf{x}_k - \mathbf{P}_{\mathcal{C}\{\mathbf{q}_1, \dots, \mathbf{q}_{k-1}\}} \mathbf{x}_k = \mathbf{x}_k - \sum_{j=1}^{k-1} \langle \mathbf{x}_k, \mathbf{q}_j \rangle \mathbf{q}_j$$

$$\mathbf{q}_k = \mathbf{v}_k / \|\mathbf{v}_k\|_2$$

- For $j = 1, \dots, p$, $\mathcal{C}\{\mathbf{x}_1, \dots, \mathbf{x}_j\} = \mathcal{C}\{\mathbf{q}_1, \dots, \mathbf{q}_j\}$ and $\mathbf{q}_j \perp \mathcal{C}\{\mathbf{x}_1, \dots, \mathbf{x}_{j-1}\}$.
- Collectively, we have $\mathbf{X} = \mathbf{Q}\mathbf{R}$ (skinny QR), where
 - $\mathbf{Q} \in \mathbb{R}^{n \times p}$ has orthonormal columns \mathbf{q}_k and thus $\mathbf{Q}^T \mathbf{Q} = \mathbf{I}_p$.
 - What's \mathbf{R} ? $\mathbf{R} = \mathbf{Q}^T \mathbf{X} \in \mathbb{R}^{p \times p}$ has entries $r_{jk} = \langle \mathbf{q}_j, \mathbf{x}_k \rangle$, which are available from the algorithm. Note $r_{jk} = 0$ for $j > k$. Thus \mathbf{R} is upper triangular.
- \mathbf{X} is over-written by \mathbf{Q} and \mathbf{R} is stored in a separate array.
- The regular Gram-Schmidt is *unstable* (we loose orthogonality due to roundoff errors) when columns of \mathbf{X} are collinear.
- *Modified Gram-Schmidt* (MGS): after each normalization step of \mathbf{v}_k , we replace $\tilde{\mathbf{x}}_j, j > k$, by its residual.

- Why MGS is better than GS? Read <http://cavern.uark.edu/~arnold/4353/CGSMGS.pdf>
- Computational cost of GS and MGS is $\sum_{k=1}^p 2n(k-1) \approx np^2$.

8 Lecture 8, Sep 23

Announcements

- HW2 due today. Submit both hardcopy and R code (LastFirstHW2.R or LastFirstHW2.Rmd)
- HW3 due next Tue.

Last time

- Cholesky decomposition with symmetric pivoting: $PAP^T = LL^T$
- Linear regression by Cholesky
- QR and linear regression
- QR by (modified) Gram-Schmidt: np^2 flops to get $X = Q_1R_1$ (thin QR)

Today

- QR by Householder
- QR by Givens

QR by Householder



Photograph by [Paul Halmos](#)

Assume $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_p) \in \mathbb{R}^{n \times p}$ has full column rank.

- Idea: $\mathbf{H}_p \cdots \mathbf{H}_2 \mathbf{H}_1 \mathbf{X} = \begin{pmatrix} \mathbf{R}_1 \\ \mathbf{0} \end{pmatrix}$, where $\mathbf{H}_j \in \mathbb{R}^{n \times n}$ are the Householder transformation matrix. It yields the full QR where $\mathbf{Q} = \mathbf{H}_1 \cdots \mathbf{H}_p \in \mathbb{R}^{n \times n}$. Recall GS/MGS only produces the thin QR decomposition.
- For arbitrary $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ with $\|\mathbf{v}\|_2 = \|\mathbf{w}\|_2$, we can construct a *Householder matrix* $\mathbf{H} = \mathbf{I}_n - 2\mathbf{u}\mathbf{u}^\top$, $\mathbf{u} = -\frac{1}{\|\mathbf{v}-\mathbf{w}\|_2}(\mathbf{v} - \mathbf{w})$, that carries \mathbf{v} to \mathbf{w} :

$$\mathbf{H}\mathbf{v} = \mathbf{w}.$$

\mathbf{H} is symmetric and orthogonal. Calculation of Householder vector \mathbf{u} costs $2n$ flops.




FIG. 2.3.1. Reflection of a vector \mathbf{a} in a hyperplane with normal \mathbf{u} .

- Now choose \mathbf{H}_1 to zero the first column of \mathbf{X} below diagonal

$$\mathbf{H}_1 \mathbf{x}_1 = \begin{pmatrix} \|\mathbf{x}_1\|_2 \\ 0 \\ \vdots \\ 0 \end{pmatrix}.$$

Take \mathbf{H}_2 to zero the second column below diagonal; ...

$$\mathbf{H}_2 \mathbf{H}_1 \mathbf{A} = \begin{bmatrix} \times & \times & \times & \times & \times \\ 0 & \times & \times & \times & \times \\ 0 & 0 & \boxtimes & \times & \times \\ 0 & 0 & \boxtimes & \times & \times \\ 0 & 0 & \boxtimes & \times & \times \\ 0 & 0 & \boxtimes & \times & \times \end{bmatrix}$$

In general, choose the j -th Householder transform $\mathbf{H}_j = \mathbf{I}_n - 2\mathbf{u}_j\mathbf{u}_j^\top$, where $\mathbf{u}_j = \begin{pmatrix} \mathbf{0}_{j-1} \\ \tilde{\mathbf{u}}_j \end{pmatrix}$, $\tilde{\mathbf{u}}_j \in \mathbb{R}^{n-j+1}$, to zero the j -th column below diagonal. \mathbf{H}_j takes the form

$$\mathbf{H}_j = \begin{pmatrix} \mathbf{I}_{j-1} & \\ & \mathbf{I}_{n-j+1} - 2\tilde{\mathbf{u}}_j\tilde{\mathbf{u}}_j^\top \end{pmatrix} = \begin{pmatrix} \mathbf{I}_{j-1} & \\ & \tilde{\mathbf{H}}_j \end{pmatrix}.$$

- Applying a Householder transform $\mathbf{H} = \mathbf{I} - 2\mathbf{u}\mathbf{u}^\top$ to a matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$

$$\mathbf{H}\mathbf{X} = \mathbf{X} - 2\mathbf{u}(\mathbf{u}^\top \mathbf{X})$$

costs $2np$ flops. *We never explicitly form the Householder matrices.*

Note applying \mathbf{H}_j to \mathbf{X} only needs $2(n-j+1)(p-j+1)$ flops.

- QR by Householder: $\mathbf{H}_p \cdots \mathbf{H}_1 \mathbf{X} = \begin{pmatrix} \mathbf{R}_1 \\ \mathbf{0} \end{pmatrix}$.
- The process is done in place. Upper triangular part of \mathbf{X} is overwritten by \mathbf{R}_1 and the essential Householder vectors ($\tilde{\mathbf{u}}_{j1}$ is normalized to 1) are stored in $\mathbf{X}(j:n, j)$.
- At j -th stage
 1. computing the Householder vector $\tilde{\mathbf{u}}_j$ costs $2(n-j+1)$ flops
 2. applying the Householder transform $\tilde{\mathbf{H}}_j$ to the $\mathbf{X}(j:n, j:p)$ block costs $2(n-j+1)(p-j+1)$ flops

In total we need $\sum_{j=1}^p [2(n-j+1) + 2(n-j+1)(p-j+1)] \approx np^2 - \frac{1}{3}p^3$ flops.

- Where is \mathbf{Q} ? $\mathbf{Q} = \mathbf{H}_1 \cdots \mathbf{H}_p$.
In some applications, it's necessary to form the orthogonal matrix \mathbf{Q} .
Accumulating \mathbf{Q} costs another $np^2 - \frac{1}{3}p^3$ flops.
- When computing $\mathbf{Q}^T \mathbf{v}$ or $\mathbf{Q} \mathbf{v}$ as in some applications (e.g., solve linear equation using QR), no need to form \mathbf{Q} . Simply apply Householder transforms successively.
`qr.qy()` and `qr.qty()` in R do this.
- Computational cost of Householder QR for linear regression: $np^2 - \frac{1}{3}p^3$ (regression coefficients and $\hat{\sigma}^2$) or more (fitted values, s.e., ...).

Rank deficient \mathbf{X} : Householder QR with column pivoting

$\mathbf{X} \in \mathbb{R}^{n \times p}$ may not have full column rank.

- Idea (due to Businger and Golub 1965): at the j -th stage, swap the column in $\mathbf{X}(j : n, j : p)$ with maximum ℓ_2 norm to be the pivot column. If the maximum ℓ_2 norm is 0, it stops, ending with

$$\mathbf{X}\mathbf{\Pi} = \mathbf{Q} \begin{pmatrix} \mathbf{R}_{11} & \mathbf{R}_{12} \\ \mathbf{0}_{(n-r) \times r} & \mathbf{0}_{(n-r) \times (p-r)} \end{pmatrix},$$

where $\mathbf{\Pi} \in \mathbb{R}^{p \times p}$ is a permutation matrix and r is the rank of \mathbf{X} . QR with column pivoting is rank revealing.

- The overhead of re-computing the column norms can be reduced by the property

$$\mathbf{Q} \mathbf{z} = \begin{pmatrix} \alpha \\ \boldsymbol{\omega} \end{pmatrix} \Rightarrow \|\boldsymbol{\omega}\|_2^2 = \|\mathbf{z}\|_2^2 - \alpha^2$$

for any orthogonal matrix \mathbf{Q} .

- In R, the `qr()` function is a wrapper for various LINPACK (default) and LAPACK routines. It performs Householder QR with column pivoting and returns
 - `$qr`: a matrix of same size as input matrix
 - `$rank`: rank of the input matrix

- `$pivot`: pivot vector
- `$aux`: normalizing constants of Householder vectors

Auxiliary functions `qr.coef()`, `qr.resid()`, `qr.qy()`, `qr.qty()`, `qr.solve()`, ... are very helpful.

QR by Givens rotation (JM 5.7-5.8)



- Householder transform \mathbf{H}_j introduces batch of zeros into a vector.
- Givens transform (aka Givens rotation, Jacobi rotation, plane rotation) selectively zeros one element of a vector.
- Overall QR by Givens rotation is less efficient than the Householder method, but is better suited for matrices with structured patterns of nonzero elements.
- Givens/Jacobi rotations:

$$\mathbf{G}(i, k, \theta) = \begin{pmatrix} 1 & & 0 & & 0 & & 0 \\ \vdots & \ddots & \vdots & & \vdots & & \vdots \\ 0 & & c & & s & & 0 \\ \vdots & & \vdots & \ddots & \vdots & & \vdots \\ 0 & & -s & & c & & 0 \\ \vdots & & \vdots & & \vdots & \ddots & \vdots \\ 0 & & 0 & & 0 & & 1 \end{pmatrix},$$

where $c = \cos(\theta)$ and $s = \sin(\theta)$. $\mathbf{G}(i, k, \theta)$ is orthogonal.

- Pre-multiplication by $\mathbf{G}(i, k, \theta)^T$ rotates counterclockwise θ radians in the (i, k) coordinate plane. If $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{y} = \mathbf{G}(i, k, \theta)^T \mathbf{x}$, then

$$y_j = \begin{cases} cx_i - sx_k & j = i \\ sx_i + cx_k & j = k \\ x_j & j \neq i, k \end{cases} .$$

Apparently if we choose $\tan(\theta) = -x_k/x_i$, or equivalently,

$$c = \frac{x_i}{\sqrt{x_i^2 + x_k^2}}, \quad s = \frac{-x_k}{\sqrt{x_i^2 + x_k^2}},$$

then $y_k = 0$.

- Pre-applying Givens transform $\mathbf{G}(i, k, \theta)^T \in \mathbb{R}^{n \times n}$ to a matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$ only effects two rows of \mathbf{A} :

$$\mathbf{A}([i, k], :) \leftarrow \begin{pmatrix} c & s \\ -s & c \end{pmatrix}^T \mathbf{A}([i, k], :),$$

costing $2m$ flops.

- Post-applying Givens transform $\mathbf{G}(i, k, \theta) \in \mathbb{R}^{m \times m}$ to a matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$ only effects two columns of \mathbf{A} :

$$\mathbf{A}(:, [i, k]) \leftarrow \mathbf{A}(:, [i, k]) \begin{pmatrix} c & s \\ -s & c \end{pmatrix},$$

costing $2n$ flops.

- QR by Givens: $\mathbf{G}_t^T \cdots \mathbf{G}_1^T \mathbf{X} = \begin{pmatrix} \mathbf{R}_1 \\ \mathbf{0} \end{pmatrix}$

$$\begin{array}{c} \begin{bmatrix} \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \end{bmatrix} \xrightarrow{\langle 3,4 \rangle} \begin{bmatrix} \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \mathbf{0} & \times & \times \end{bmatrix} \xrightarrow{\langle 2,3 \rangle} \begin{bmatrix} \times & \times & \times \\ \times & \times & \times \\ \mathbf{0} & \times & \times \\ \mathbf{0} & \times & \times \end{bmatrix} \xrightarrow{\langle 1,2 \rangle} \\ \\ \begin{bmatrix} \times & \times & \times \\ \mathbf{0} & \times & \times \\ \mathbf{0} & \times & \times \\ \mathbf{0} & \times & \times \end{bmatrix} \xrightarrow{\langle 3,4 \rangle} \begin{bmatrix} \times & \times & \times \\ \mathbf{0} & \times & \times \\ \mathbf{0} & \times & \times \\ \mathbf{0} & \mathbf{0} & \times \end{bmatrix} \xrightarrow{\langle 2,3 \rangle} \begin{bmatrix} \times & \times & \times \\ \mathbf{0} & \times & \times \\ \mathbf{0} & \mathbf{0} & \times \\ \mathbf{0} & \mathbf{0} & \times \end{bmatrix} \xrightarrow{\langle 3,4 \rangle} \mathbf{R} \end{array}$$

- Zeros in \mathbf{X} can also be introduced row-by-row.
- If $\mathbf{X} \in \mathbb{R}^{n \times p}$, the total cost is $(3/2)np^2 - p^3/2$ flops and $O(np)$ square roots.
- Note each Givens transform can be summarized by a single number, which is stored in the zeroed entry of \mathbf{X} .
- Fast Givens transform avoids taking square roots.

9 Lecture 9, Sep 25

Announcements

- HW2 returned. Feedback:
 - Only Q2, Q4, and Q5 are graded. Maximum points 60.
 - Code style penalty doubles.
 - Late penalty: 5 pts/day.
 - Sketch of solution: <http://hua-zhou.github.io/teaching/st758-2014fall/hw02sol.html>. Please compare to your code carefully and understand why.
- HW3 due next Tue. FAQs at <http://hua-zhou.github.io/teaching/st758-2014fall/st758fall2014/2014/09/24/hw3-hints.html>
- HW4 posted. Due Oct 14.
- Answer to Susheela's questions:
 - *Backward accumulation algorithm* for accumulating $\mathbf{Q} = \mathbf{H}_1 \cdots \mathbf{H}_p$ from Householder vectors stored in the Household QR output costs about $2(n^2p - np^2 + p^3/3)$ flops. See Golub and Van Loan (1996, p213).
 - Computing $\mathbf{Q}^T \mathbf{Y}$, $\mathbf{Y} \in \mathbb{R}^{n \times r}$ (r right hand sides), costs about $(2np - p^2)r$ flops. See Golub and Van Loan (1996, p212).
- Answer to Meng's question: Why (how) normalize Householder vector such that the 1st element is 1? See Golub and Van Loan (1996, Algorithm 5.1.1, p212).

Last time

- QR by Householder
- QR by Givens

Today

- Sweep operator (KL 7.4-7.6, JM 5.12)

Sweep operator

Assume $\mathbf{A} \succeq \mathbf{0}_{n \times n}$.

- KL 7.4-7.6; JM 5.12; Also see “*A tutorial on the SWEEP operator*” by James H. Goodnight. <http://www.jstor.org/stable/2683825>
- Note the (anti-symmetric) version in JM is different from the (symmetric) version in KL. I follow the convention in KL.
- *Sweep* on the k -th diagonal entry $a_{kk} \neq 0$ yields $\hat{\mathbf{A}}$ with entries

$$\begin{aligned}\hat{a}_{kk} &= -\frac{1}{a_{kk}} \\ \hat{a}_{ik} &= \frac{a_{ik}}{a_{kk}} \\ \hat{a}_{kj} &= \frac{a_{kj}}{a_{kk}} \\ \hat{a}_{ij} &= a_{ij} - \frac{a_{ik}a_{kj}}{a_{kk}}, \quad i \neq k, j \neq k.\end{aligned}$$

$n^2/2$ flops (taking into account of symmetry).

- *Inverse sweep* sends \mathbf{A} to $\check{\mathbf{A}}$ with entries

$$\begin{aligned}\check{a}_{kk} &= -\frac{1}{a_{kk}} \\ \check{a}_{ik} &= -\frac{a_{ik}}{a_{kk}} \\ \check{a}_{kj} &= -\frac{a_{kj}}{a_{kk}} \\ \check{a}_{ij} &= a_{ij} - \frac{a_{ik}a_{kj}}{a_{kk}}, \quad i \neq k, j \neq k.\end{aligned}$$

$n^2/2$ flops (taking into account of symmetry).

- $\check{\check{\mathbf{A}}} = \mathbf{A}$.
- Block form of sweep: Let the symmetric matrix \mathbf{A} be partitioned as $\mathbf{A} = \begin{pmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{pmatrix}$. If possible, sweep on the diagonal entries of \mathbf{A}_{11} yields

$$\hat{\mathbf{A}} = \begin{pmatrix} -\mathbf{A}_{11}^{-1} & \mathbf{A}_{11}^{-1} \mathbf{A}_{12} \\ \mathbf{A}_{21} \mathbf{A}_{11}^{-1} & \mathbf{A}_{22} - \mathbf{A}_{21} \mathbf{A}_{11}^{-1} \mathbf{A}_{12} \end{pmatrix}.$$

Order dose *not* matter.

- Pd and determinant: \mathbf{A} is pd if and only if each diagonal entry can be swept in succession and is positive until it is swept. When a diagonal entry of a pd matrix \mathbf{A} is swept, it becomes negative and remains negative thereafter. Taking the product of diagonal entries just before each is swept yields the determinant of \mathbf{A} .

- Linear regression by sweep. Sweep on $\begin{pmatrix} \mathbf{X}^\top \mathbf{X} & \mathbf{X}^\top \mathbf{y} \\ \mathbf{y}^\top \mathbf{X} & \mathbf{y}^\top \mathbf{y} \end{pmatrix}$ yields

$$\begin{pmatrix} -(\mathbf{X}^\top \mathbf{X})^{-1} & (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y} \\ \mathbf{y}^\top \mathbf{X} (\mathbf{X} \mathbf{X})^{-1} & \mathbf{y}^\top \mathbf{y} - \mathbf{y}^\top \mathbf{X} (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y} \end{pmatrix} = \begin{pmatrix} -\frac{1}{\sigma^2} \text{Var}(\hat{\boldsymbol{\beta}}) & \hat{\boldsymbol{\beta}} \\ \hat{\boldsymbol{\beta}}^\top & \|\mathbf{y} - \hat{\mathbf{y}}\|_2^2 \end{pmatrix}.$$

In total $np^2/2 + p^3/2$ flops.

- Sweep is useful for stepwise regression, (conditional) multivariate normal density calculation, MANOVA, ...
- Warning: the `sweep()` function in R has nothing do to with the sweep operator here.
- Demo code: <http://hua-zhou.github.io/teaching/st758-2014fall/sweep.html>

10 Lecture 10, Sep 30

Announcements

- HW3 due today. Submit hard copy + code (LastFirstHW3.R or LastFirstHW3.Rmd)
- HW4 posted. Due Oct 14.
- TA office hours on Oct 8 @ 10AM-12PM?
- Some FAQs:
 - (Maggie) If $\mathbf{A} \in \mathbb{R}^{n \times p}$ has full rank, Householder QR gives $\mathbf{A} = \mathbf{Q}_1 \mathbf{R}_1$ and MGS gives $\mathbf{A} = \mathbf{Q}_2 \mathbf{R}_2$, where $\mathbf{Q}_1, \mathbf{Q}_2 \in \mathbb{R}^{n \times p}$, $\mathbf{Q}_1^T \mathbf{Q}_1 = \mathbf{Q}_2^T \mathbf{Q}_2 = \mathbf{I}_p$, and $\mathbf{R}_1, \mathbf{R}_2 \in \mathbb{R}^{p \times p}$ are upper triangular with positive diagonal entries. How do we know $\mathbf{Q}_1 = \mathbf{Q}_2$?
 - (Meng) Block sweep simply means sweeping diagonal entries in that block sequentially. Order does *not* matter.
 - Check “80 character rule” in R Studio.
 - (Liuyi) What exactly is that `pivot` from the output of `qr()` and `chol()` (with `pivot = TRUE`) functions in R? See hints at <http://hua-zhou.github.io/teaching/st758-2014fall/st758fall12014/2014/09/24/hw3-hints.html>

Last time

- HW2 review
- Sweep operator (KL 7.4-7.6, JM 5.12)

Today

- Summary of numerical methods for linear regression
- Summary of solving linear equations: overdetermined system
- Condition number for solving linear equations

Summary of linear regression: Table on KL p105

Method	Flops	Remarks	Software	Stability
Sweep	$np^2/2 + p^3/2$	$(\mathbf{X}^\top \mathbf{X})^{-1}$ available	SAS	less stable
Cholesky	$np^2/2 + p^3/6$			less stable
QR by Householder	$np^2 - p^3/3$		R	
QR by MGS	np^2	\mathbf{Q}_1 available		more stable

Table 1: Numerical methods for linear regression. In order of stability.

Remarks:

- When $n \gg p$, sweep and Cholesky are twice faster than QR and need less space. But QR methods are more stable and produce numerically more accurate solution.
- Although sweep is slower than Cholesky, it yields standard errors and so on.
- Sweep is useful for stepwise regression, multivariate normal calculation, and numerous other statistical applications.
- MGS appears slower than Householder, but it yields \mathbf{Q}_1 .

“There is simply no such thing as a universal ‘gold standard’ when it comes to algorithms”.

Summary of solving linear equations

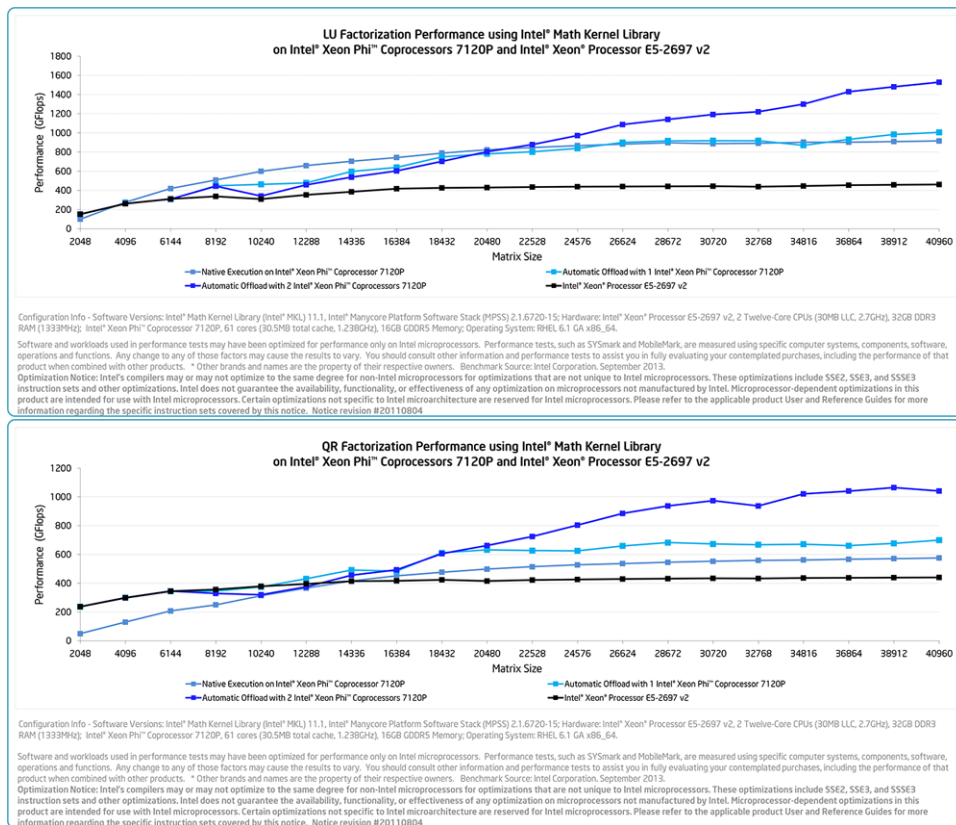
Consider linear system $\mathbf{Ax} = \mathbf{b}$.

- We now know some good numerical methods for the least squares problem, which is essentially “solving” an overdetermined system (a tall \mathbf{A}).
- Table 2 compares the flops of some methods (in order of stability) for solving a square (unstructured) $\mathbf{A} \in \mathbb{R}^{n \times n}$.

Method	Flops	Stability
Gaussian elimination	$n^3/3$	less stable
QR by Householder	$(2/3)n^3$	
QR by MGS	n^3	
SVD	$6n^3$	most stable

Table 2: Flops of different numerical methods for $n \times n$ square linear systems, assuming availability of the right hand side at time of decomposition.

- Flop count is not everything. GE/LU has a higher memory traffic and vectorization overheads, and QR approach is comparable in efficiency. QR methods are more stable.



- Solve an underdetermined system (a flat $A \in \mathbb{R}^{m \times n}$ of full row rank) by QR –

version 1. First compute QR on \mathbf{A}^T

$$\mathbf{A}^T = \mathbf{Q}\mathbf{R} = \mathbf{Q} \begin{pmatrix} \mathbf{R}_1 \\ \mathbf{0}_{(n-m) \times m} \end{pmatrix}.$$

Then $\mathbf{A}\mathbf{x} = \mathbf{b}$ becomes

$$(\mathbf{Q}\mathbf{R})^T \mathbf{x} = \begin{pmatrix} \mathbf{R}_1^T & \mathbf{0}_{m \times (n-m)} \end{pmatrix} \begin{pmatrix} \mathbf{z}_1 \\ \mathbf{z}_2 \end{pmatrix} = \mathbf{b},$$

where

$$\mathbf{Q}^T \mathbf{x} = \begin{pmatrix} \mathbf{z}_1 \\ \mathbf{z}_2 \end{pmatrix}.$$

\mathbf{z}_1 is determined from $\mathbf{R}_1^T \mathbf{z}_1 = \mathbf{b}$. If we take $\mathbf{z}_2 = \mathbf{0}_{n-m}$, then we obtain the minimum norm solution (why?).

- Solve an underdetermined system (a flat $\mathbf{A} \in \mathbb{R}^{m \times n}$ of full row rank) by QR – version 2. First compute QR with column pivoting on \mathbf{A}

$$\mathbf{A}\mathbf{\Pi} = \mathbf{Q} \begin{pmatrix} \mathbf{R}_1 & \mathbf{R}_2 \end{pmatrix},$$

where $\mathbf{R}_1 \in \mathbb{R}^{m \times m}$ is upper triangular and $\mathbf{R}_2 \in \mathbb{R}^{m \times (n-m)}$. Thus $\mathbf{A}\mathbf{x} = \mathbf{b}$ transforms to

$$\begin{pmatrix} \mathbf{R}_1 & \mathbf{R}_2 \end{pmatrix} \begin{pmatrix} \mathbf{z}_1 \\ \mathbf{z}_2 \end{pmatrix} = \mathbf{Q}^T \mathbf{b},$$

where

$$\mathbf{\Pi}^T \mathbf{x} = \begin{pmatrix} \mathbf{z}_1 \\ \mathbf{z}_2 \end{pmatrix}.$$

One solution is obtained by $\mathbf{z}_1 = \mathbf{R}_1^{-1} \mathbf{Q}^T \mathbf{b}$ and $\mathbf{z}_2 = \mathbf{0}_{n-m}$. It is not guaranteed to be of minimum norm.

Condition number for linear equations (matrix inversion)

- Assume $\mathbf{A} \in \mathbb{R}^{n \times n}$ is nonsingular and consider the system of linear equation $\mathbf{A}\mathbf{x} = \mathbf{b}$. The solution is $\mathbf{x} = \mathbf{A}^{-1} \mathbf{b}$. We want to know how the solution changes with a small perturbation of the input \mathbf{b} (or \mathbf{A}).

- Let $\tilde{\mathbf{b}} = \mathbf{b} + \Delta\mathbf{b}$. Then $\tilde{\mathbf{x}} = \mathbf{A}^{-1}(\mathbf{b} + \Delta\mathbf{b}) = \mathbf{x} + \Delta\mathbf{x}$. Thus

$$\|\Delta\mathbf{x}\| = \|\mathbf{A}^{-1}\Delta\mathbf{b}\| \leq \|\mathbf{A}^{-1}\| \|\Delta\mathbf{b}\|.$$

Because $\mathbf{b} = \mathbf{A}\mathbf{x}$, $\frac{1}{\|\mathbf{x}\|} \leq \|\mathbf{A}\| \frac{1}{\|\mathbf{b}\|}$. This results

$$\frac{\|\Delta\mathbf{x}\|}{\|\mathbf{x}\|} \leq \|\mathbf{A}\| \|\mathbf{A}^{-1}\| \frac{\|\Delta\mathbf{b}\|}{\|\mathbf{b}\|}.$$

- $\kappa(\mathbf{A}) = \|\mathbf{A}\| \|\mathbf{A}^{-1}\|$ is called the *condition number for inversion*. It depends on the matrix norm being used. κ_p means condition number defined from matrix- p norm.
- Large condition number means “bad”.
- Some useful facts

$$\begin{aligned} \kappa(\mathbf{A}) &= \kappa(\mathbf{A}^{-1}) \\ \kappa(c\mathbf{A}) &= \kappa(\mathbf{A}) \\ \kappa(\mathbf{A}) &\geq 1 \\ \kappa_1(\mathbf{A}) &= \kappa_\infty(\mathbf{A}^\top) \\ \kappa_2(\mathbf{A}) &= \kappa_2(\mathbf{A}^\top) = \frac{\sigma_1(\mathbf{A})}{\sigma_n(\mathbf{A})} \\ \kappa_2(\mathbf{A}^\top\mathbf{A}) &= \frac{\lambda_1(\mathbf{A}^\top\mathbf{A})}{\lambda_n(\mathbf{A}^\top\mathbf{A})} = \kappa_2^2(\mathbf{A}) \geq \kappa_2(\mathbf{A}). \end{aligned}$$

The last fact says that the condition number of $\mathbf{A}^\top\mathbf{A}$ can be much larger than that of \mathbf{A} .

- The smallest singular value σ_n indicates the “distance to the trouble”.
- Condition number for the least squares problem is more complicated. Roughly speaking, the method based on normal equation (Cholesky, sweep) has a condition depending on $\kappa_2(\mathbf{X})^2$. QR for a “small residuals” least squares problem has a condition depending on $\kappa_2(\mathbf{X})$.
- Numerically, consider the simple case

$$\mathbf{X} = \begin{pmatrix} 1 & 1 \\ 10^{-3} & 0 \\ 0 & 10^{-3} \end{pmatrix}.$$

Forming normal equation yields a singular Gramian matrix

$$\mathbf{X}^T \mathbf{X} = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$$

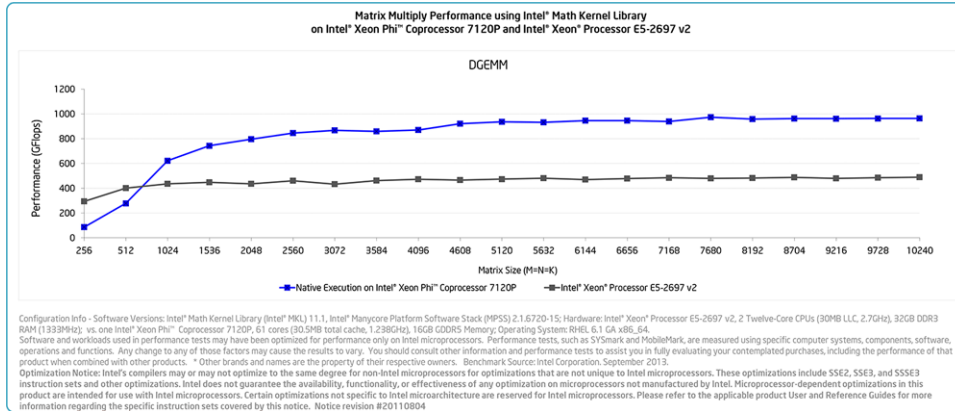

if executed with a precision of 6 digits.

- In R, the `kappa()` function (wrapper of `DTRCON` in LAPACK and `DTRCO` in LINPACK) computes or approximates (default) the condition number of a matrix.
- In regression problems, standardizing the predictors could improve the condition. See demo on the Longley data <http://hua-zhou.github.io/teaching/st758-2014fall/longleycond.html>.
- In design of experiments (DoE), people favor orthogonal design. Why?


11 Lecture 11, Oct 2

Announcements

- HW3 returned. Feedback:
 - Only Q2(d) and Q3 are graded. Total points is 50.
 - Style issues: penalty doubles.
 - Late penalty: 5 pts/day.
 - Q1, probabilistic proof?
 - By default, `qr()` function does Householder QR *with column pivoting!*
Need to use the `pivot` vector to permute results.
 - Make good use of `crossprod()` and `tcrossprod()` functions to compute Gramian matrix, avoid unnecessary matrix transpose, and so on. Check source code:
`/R-3.1.1/src/main/names.c`
`/R-3.1.1/src/main/array.c`
 - How to generate rank deficient matrix?
 - Efficient implementation of sweep operator?
 - Sketch of solution: <http://hua-zhou.github.io/teaching/st758-2014fall/hw03sol.html>. Please compare to your code carefully and understand why.
- HW4 due Oct 14.
- TA office hours on Wed Oct 8. No TA office hours on Fri Oct 10 (fall break).
- Caleb's question: Nvidia Tesla vs Intel Xeon Phi



Our department has at least 2 servers each with 4 Nvidia Tesla M2070Q GPUs. Each Tesla M2070Q has 6G memory (5.25G with ECC), 786K L2 cache, 448 cores @ 1.15GHz, and theoretical throughput of 1288 SP GFLOPS or 512 DP GFLOPS.



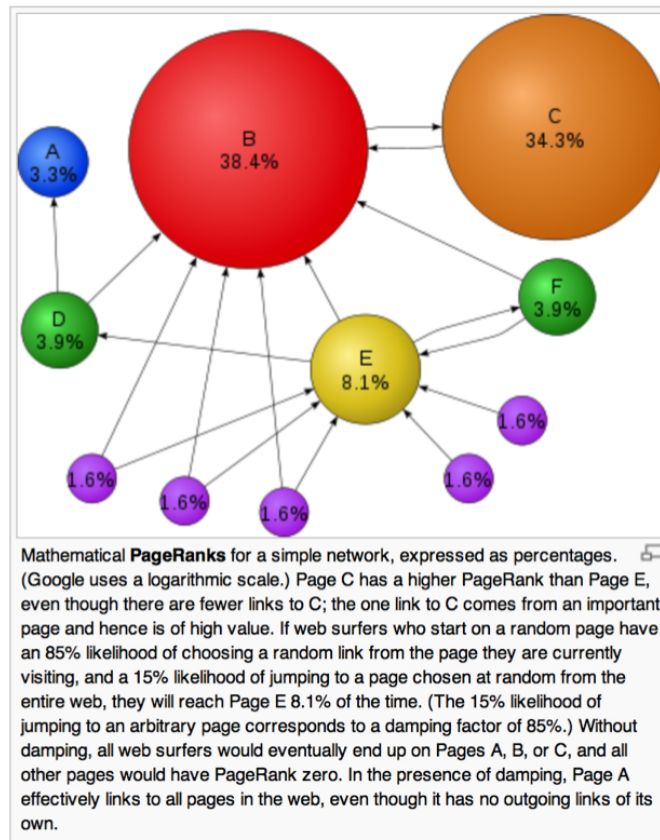
Last time

- Summary of linear regressions
- Summary of numerical methods for solving linear equations: tall \mathbf{A} (least squares, over-determined system), square \mathbf{A} , and flat \mathbf{A} (underdetermined system)
- Condition number for solving linear equations

Today

- Iterative solvers for linear systems

Iterative method for solving linear equations: introduction



- Direct method (flops fixed *a priori*) vs iterative methods:

- Direct method (GE/LU, Cholesky, QR, SVD): good for dense, small or moderate sized, unstructured \mathbf{A}
- Iterative methods (Jacobi, Gauss-Seidal, SOR, conjugate-gradient, GMRES): good for large, sparse, or structured linear system, parallel computing, warm start
- PageRank (HW4):
 - $\mathbf{A} \in \{0, 1\}^{n \times n}$ the connectivity matrix with entries

$$a_{ij} = \begin{cases} 1 & \text{if page } i \text{ links to page } j \\ 0 & \text{otherwise} \end{cases}.$$

$n \approx 10^9$ in Sep 2014.


- $r_i = \sum_j a_{ij}$ is the out-degree of page i .
- Imagine a random surfer wandering on internet according to following rules:
 - * From a page i with $r_i > 0$
 - with probability p , (s)he randomly chooses a link on page i (uniformly) and follows that link to the next page
 - with probability $1 - p$, (s)he randomly chooses one page from the set of all n pages (uniformly) and proceeds to that page
 - * From a page i with $r_i = 0$ (a dangling page), (s)he randomly chooses one page from the set of all n pages (uniformly) and proceeds to that page

The process defines a Markov chain on the space of n pages.

- Stationary distribution of this Markov chain gives the ranks (probabilities) of each page.
- Stationary distribution is the top left eigenvector of the transition matrix \mathbf{P} corresponding to eigenvalue 1. Equivalently it can be cast as a linear equation.
- Largest matrix computation in world (?).

- GE/LU will take $(10^9)^3/3/10^{12} \approx 3.33 \times 10^{14}$ seconds $\approx 1 \times 10^7$ years on a tera-flop supercomputer!
- Iterative methods come to the rescue.

Jacobi method



Solve linear system $\mathbf{Ax} = \mathbf{b}$.

- Jacobi iteration:

$$x_i^{(t+1)} = \frac{b_i - \sum_{j=1}^{i-1} a_{ij}x_j^{(t)} - \sum_{j=i+1}^n a_{ij}x_j^{(t)}}{a_{ii}}.$$

- Requires non-zero diagonal element!
- One round costs n^2 flops with an unstructured \mathbf{A} . Gain over GE/LU if converges in $o(n)$ iterations. Saving is huge for *sparse* or *structured* \mathbf{A} . By structured, we mean the matrix-vector multiplication \mathbf{Av} is fast.
- Splitting: $\mathbf{A} = \mathbf{L} + \mathbf{D} + \mathbf{U}$.
- Jacobi: $\mathbf{D}\mathbf{x}^{(t+1)} = -(\mathbf{L} + \mathbf{U})\mathbf{x}^{(t)} + \mathbf{b}$, i.e.,

$$\mathbf{x}^{(t+1)} = -\mathbf{D}^{-1}(\mathbf{L} + \mathbf{U})\mathbf{x}^{(t)} + \mathbf{D}^{-1}\mathbf{b}.$$

Gauss-Seidel



- Gauss-Seidel iteration:

$$x_i^{(t+1)} = \frac{b_i - \sum_{j=1}^{i-1} a_{ij}x_j^{(t+1)} - \sum_{j=i+1}^n a_{ij}x_j^{(t)}}{a_{ii}}.$$

- With splitting, $(\mathbf{D} + \mathbf{L})\mathbf{x}^{(t+1)} = -\mathbf{U}\mathbf{x}^{(t)} + \mathbf{b}$, i.e.,

$$\mathbf{x}^{(t+1)} = -(\mathbf{D} + \mathbf{L})^{-1}\mathbf{U}\mathbf{x}^{(t)} + (\mathbf{D} + \mathbf{L})^{-1}\mathbf{b}.$$

- GS converges for any $\mathbf{x}^{(0)}$ for symmetric and pd \mathbf{A} .
- Convergence rate of Gauss-Seidel is the spectral radius of the $(\mathbf{D} + \mathbf{L})^{-1}\mathbf{U}$.
- Comparing Jacobi and GS, Jacobi is particularly attractive for parallel computing.

Successive over-relaxation (SOR)

- SOR: $x_i^{(t+1)} = \omega(b_i - \sum_{j=1}^{i-1} a_{ij}x_j^{(t+1)} - \sum_{j=i+1}^n a_{ij}x_j^{(t)})/a_{ii} + (1 - \omega)x_i^{(t)}$, i.e.,

$$\mathbf{x}^{(t+1)} = (\mathbf{D} + \omega\mathbf{L})^{-1}[(1 - \omega)\mathbf{D} - \omega\mathbf{U}]\mathbf{x}^{(t)} + (\mathbf{D} + \omega\mathbf{L})^{-1}(\mathbf{D} + \mathbf{L})^{-1}\omega\mathbf{b}.$$

- Need to pick $\omega \in [0, 1]$ beforehand, with the goal of improving convergence rate.

Conjugate gradient method

Solving $\mathbf{Ax} = \mathbf{b}$ is equivalent to minimizing the quadratic function $\frac{1}{2}\mathbf{x}^T \mathbf{Ax} - \mathbf{b}^T \mathbf{x}$. To do later, when we study optimization. Conjugate gradient and its variants are the top-notch iterative methods for solving huge, structured linear systems.

Table 1. Kershaw's results for a fusion problem.

Method	Number of iterations
Gauss Seidel	208,000
Block successive overrelaxation methods	765
Incomplete Cholesky conjugate gradients	25

12 Lecture 12, Oct 7

Announcements

- HW4 due next Tue Oct 14
- HW5 will be posted this week
- TA office hours this week: Wed Oct 8 @ 10A-12P
- Dr. Zhou office hours this week: Tue Oct 7 @ 4P-5P, Thu Oct 9 by appointment

Last time

- HW3 review
- Iterative methods for solving linear equations

Today

- A catalog of “easy” linear systems – last topic in solving linear equations
- Eigen-decomposition and SVD

A list of “easy” linear systems

Consider $\mathbf{Ax} = \mathbf{b}$, $\mathbf{A} \in \mathbb{R}^{n \times n}$. Or, consider matrix inverse (if you want). \mathbf{A} can be huge. Keep massive data in mind: 1000 Genome Project, NetFlix, Google PageRank, finance, spatial statistics, ... We should be alert to many easy linear systems. *Don't waste computing resources by bad choices of algorithms!*

- Diagonal: n flops.
- Tridiagonal/banded: Band LU, band Cholesky, ... roughly $O(n)$ flops
- Triangular: $n^2/2$ flops
- Block diagonal: Suppose $n = \sum_i n_i$. $(\sum_i n_i)^3/3$ vs $\sum_i n_i^3/3$.
- Kronecker product: $(\mathbf{A} \otimes \mathbf{B})^{-1} = \mathbf{A}^{-1} \otimes \mathbf{B}^{-1}$, $(\mathbf{C}^\top \otimes \mathbf{A})\text{vec}\mathbf{B} = \text{vec}(\mathbf{ABC})$
fits iterative method.

- Sparsity: iterative method, or sparse matrix decomposition.

Remark: Probably the easiest recognizable structure. Familiarize yourself with the sparse matrix computation tools in MATLAB, R (Matrix package), MKL (sparse BLAS), ... as soon as possible.

- Easy plus low rank: $\mathbf{U} \in \mathbb{R}^{n \times m}$, $\mathbf{V} \in \mathbb{R}^{n \times m}$, $m \ll n$. Woodbury formula

$$(\mathbf{A} + \mathbf{U}\mathbf{V}^\top)^{-1} = \mathbf{A}^{-1} - \mathbf{A}^{-1}\mathbf{U}(\mathbf{I}_m + \mathbf{V}^\top\mathbf{A}^{-1}\mathbf{U})^{-1}\mathbf{V}^\top\mathbf{A}^{-1}.$$

Keep HW2 Q5 in mind.

- Easy plus border: For \mathbf{A} pd and \mathbf{V} full row rank,

$$\begin{pmatrix} \mathbf{A} & \mathbf{V}^\top \\ \mathbf{V} & \mathbf{0} \end{pmatrix}^{-1} = \begin{pmatrix} \mathbf{A}^{-1} - \mathbf{A}^{-1}\mathbf{V}^\top(\mathbf{V}\mathbf{A}^{-1}\mathbf{V}^\top)^{-1}\mathbf{V}\mathbf{A}^{-1} & \mathbf{A}^{-1}\mathbf{V}^\top(\mathbf{V}\mathbf{A}^{-1}\mathbf{V}^\top)^{-1} \\ (\mathbf{V}\mathbf{A}^{-1}\mathbf{V}^\top)^{-1}\mathbf{V}\mathbf{A}^{-1} & -(\mathbf{V}\mathbf{A}^{-1}\mathbf{V}^\top)^{-1} \end{pmatrix}.$$

- Orthogonal: n^2 flops *at most*. Permutation matrix, Householder matrix, Jacobi matrix, ... take less.
- Toeplitz systems:

$$\mathbf{T} = \begin{pmatrix} r_0 & r_1 & r_2 & r_3 \\ r_{-1} & r_0 & r_1 & r_2 \\ r_{-2} & r_{-1} & r_0 & r_1 \\ r_{-3} & r_{-2} & r_{-1} & r_0 \end{pmatrix}.$$

$\mathbf{T}\mathbf{x} = \mathbf{b}$, where \mathbf{T} is pd and Toeplitz, can be solved in $O(n^2)$ flops. Durbin algorithm (Yule-Walker equation), Levinson algorithm (general \mathbf{b}), Trench algorithm (inverse). These matrices occur in auto-regressive models and econometrics.

- Circulant systems: Toeplitz matrix with wraparound

$$\mathbf{C}(\mathbf{z}) = \begin{pmatrix} z_0 & z_4 & z_3 & z_2 & z_1 \\ z_1 & z_0 & z_4 & z_3 & z_2 \\ z_2 & z_1 & z_0 & z_4 & z_3 \\ z_3 & z_2 & z_1 & z_0 & z_4 \\ z_4 & z_3 & z_2 & z_1 & z_0 \end{pmatrix},$$

DCT (discrete cosine transform) and DST (discrete sine transform).
FFT type algorithms.

- Vandermonde matrix: such as in interpolation and approximation problems

$$\mathbf{V}(x_0, \dots, x_n) = \begin{pmatrix} 1 & 1 & \cdots & 1 \\ x_0 & x_1 & \cdots & x_n \\ \vdots & \vdots & & \vdots \\ x_0^n & x_1^n & \cdots & x_n^n \end{pmatrix}.$$

$\mathbf{V}\mathbf{x} = \mathbf{b}$ or $\mathbf{V}^T\mathbf{x} = \mathbf{b}$ can be solved in $O(n^2)$ flops.

- Cauchy-like matrices:

$$\mathbf{\Omega}\mathbf{A} - \mathbf{A}\mathbf{\Lambda} = \mathbf{R}\mathbf{S}^T,$$

where $\mathbf{\Omega} = \text{diag}(\omega_1, \dots, \omega_n)$ and $\mathbf{\Lambda} = (\lambda_1, \dots, \lambda_n)$. $O(n)$ flops for LU and QR.

- Structured-rank problems: semiseparable matrices (LU and QR takes $O(n)$ flops), quasiseparable matrices, ...
- Fast multiple method (FMM) for kernel matrix.
- ...

Other computations such as matrix-vector multiplication with these “easy” matrices are typically fast too.

Bottom line: Don’t blindly use `solve()`.

Linear algebra review: eigen-decomposition

Our last topic on numerical linear algebra is eigen-decomposition and singular value decomposition (SVD). We already saw the wide applications of QR decomposition in least squares problem and solving square and under-determined linear equations. Eigen-decomposition and SVD can be deemed as more thorough orthogonalization of a matrix. We start with a brief review of the related linear algebra.

- *Eigenvalues* are defined as roots of the characteristic equation $\det(\lambda\mathbf{I}_n - \mathbf{A}) = 0$.
- If λ is an eigenvalue of \mathbf{A} , then there exist non-zero $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ such that $\mathbf{A}\mathbf{x} = \lambda\mathbf{x}$ and $\mathbf{y}^T\mathbf{A} = \lambda\mathbf{y}^T$. \mathbf{x} and \mathbf{y} are called the (column) *eigenvector* and *row eigenvector* of \mathbf{A} associated with the eigenvalue λ .

- \mathbf{A} is singular if and only if it has at least one 0 eigenvalue.
- Eigenvectors associated with distinct eigenvalues are linearly independent.
- Eigenvalues of an upper or lower triangular matrix are its diagonal entries:
 $\lambda_i = a_{ii}$.
- Eigenvalues of an idempotent matrix are either 0 or 1.
- Eigenvalues of an orthogonal matrix have complex modulus 1.
- In most statistical applications, we deal with eigenvalues/eigenvectors of symmetric matrices. The eigenvalues and eigenvectors of a real *symmetric* matrix are real.
- Eigenvectors associated with distinct eigenvalues of a symmetry matrix are orthogonal.
- Eigen-decomposition of a symmetric matrix: $\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$, where
 - $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_n)$
 - columns of \mathbf{U} are the eigenvectors, which are (or can be chosen to be) mutually orthonormal. That is \mathbf{U} is an orthogonal matrix.
- A real symmetric matrix is positive semidefinite (positive definite) if and only if all eigenvalues are nonnegative (positive).
- *Spectral radius* $\rho(\mathbf{A}) = \max_i |\lambda_i|$.
- $\mathbf{A} \in \mathbb{R}^{n \times n}$ a square matrix (not required to be symmetric), then $\text{tr}(\mathbf{A}) = \sum_i \lambda_i$ and $|\mathbf{A}| = \prod_i \lambda_i$.

13 Lecture 13, Oct 14

Announcements

- HW4 due today: hardcopy + email code (LastFirstHW4.R or LastFirstHW4.Rmd)
- HW5 posted, due next Tue Oct 21

Last time

- A list of “easy” linear system
- Linear algebra review: eigen-decomposition

Today

- Linear algebra review: SVD
- Applications of eigen-decomposition and SVD
- Algorithms for eigen-decomposition and SVD

Linear algebra review: SVD

- Singular value decomposition (SVD): For a rectangular matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, let $p = \min\{m, n\}$, then we have the SVD

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top,$$

where

- $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_m) \in \mathbb{R}^{m \times m}$ is orthogonal
- $\mathbf{V} = (\mathbf{v}_1, \dots, \mathbf{v}_n) \in \mathbb{R}^{n \times n}$ is orthogonal
- $\mathbf{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_p) \in \mathbb{R}^{m \times n}$, $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_p \geq 0$.

σ_i are called the *singular values*, \mathbf{u}_i are the *left singular vectors*, and \mathbf{v}_i are the *right singular vectors*.

- Thin SVD. Assume $m \geq n$. \mathbf{A} can be factored as

$$\mathbf{A} = \mathbf{U}_n \mathbf{\Sigma}_n \mathbf{V}^T = \sum_{i=1}^n \sigma_i \mathbf{u}_i \mathbf{v}_i^T,$$

where

- $\mathbf{U}_n \in \mathbb{R}^{m \times n}$, $\mathbf{U}_n^T \mathbf{U}_n = \mathbf{I}_n$
- $\mathbf{V} \in \mathbb{R}^{n \times n}$, $\mathbf{V}^T \mathbf{V} = \mathbf{I}_n$
- $\mathbf{\Sigma}_n = \text{diag}(\sigma_1, \dots, \sigma_n) \in \mathbb{R}^{n \times n}$, $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0$

- Denote $\sigma(\mathbf{A}) = (\sigma_1, \dots, \sigma_p)^T$. Then

- $r = \text{rank}(\mathbf{A}) = \# \text{ nonzero singular values} = \|\sigma(\mathbf{A})\|_0$
- $\mathbf{A} = \mathbf{U}_r \mathbf{\Sigma}_r \mathbf{V}_r^T = \sum_{i=1}^r \sigma_i \mathbf{u}_i \mathbf{v}_i^T$
- $\|\mathbf{A}\|_F = (\sum_{i=1}^p \sigma_i^2)^{1/2} = \|\sigma(\mathbf{A})\|_2$
- $\|\mathbf{A}\|_2 = \sigma_1 = \|\sigma(\mathbf{A})\|_\infty$

- Assume $\text{rank}(\mathbf{A}) = r$ and partition $\mathbf{U} = (\mathbf{U}_r, \tilde{\mathbf{U}}_r) \in \mathbb{R}^{m \times m}$ and $\mathbf{V} = (\mathbf{V}_r, \tilde{\mathbf{V}}_r) \in \mathbb{R}^{n \times n}$, then

- $\mathcal{C}(\mathbf{A}) = \mathcal{C}(\mathbf{U}_r)$, $\mathcal{N}(\mathbf{A}^T) = \mathcal{C}(\tilde{\mathbf{U}}_r)$
- $\mathcal{N}(\mathbf{A}) = \mathcal{C}(\tilde{\mathbf{V}}_r)$, $\mathcal{C}(\mathbf{A}^T) = \mathcal{C}(\mathbf{V}_r)$
- $\mathbf{U}_r \mathbf{U}_r^T$ is the orthogonal projection onto $\mathcal{C}(\mathbf{A})$
- $\tilde{\mathbf{U}}_r \tilde{\mathbf{U}}_r^T$ is the orthogonal projection onto $\mathcal{N}(\mathbf{A}^T)$
- $\mathbf{V}_r \mathbf{V}_r^T$ is the orthogonal projection onto $\mathcal{C}(\mathbf{A}^T)$
- $\tilde{\mathbf{V}}_r \tilde{\mathbf{V}}_r^T$ is the orthogonal projection onto $\mathcal{N}(\mathbf{A})$

- Relation to eigen-decomposition. Using thin SVD,

$$\begin{aligned} \mathbf{A}^T \mathbf{A} &= \mathbf{V} \mathbf{\Sigma} \mathbf{U}^T \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T = \mathbf{V} \mathbf{\Sigma}^2 \mathbf{V}^T \\ \mathbf{A} \mathbf{A}^T &= \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T \mathbf{V} \mathbf{\Sigma} \mathbf{U}^T = \mathbf{U} \mathbf{\Sigma}^2 \mathbf{U}^T. \end{aligned}$$

- Another relation to eigen-decomposition. Using thin SVD,

$$\begin{pmatrix} \mathbf{0}_{n \times n} & \mathbf{A}^T \\ \mathbf{A} & \mathbf{0}_{m \times m} \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} \mathbf{V} & \mathbf{V} \\ \mathbf{U} & -\mathbf{U} \end{pmatrix} \begin{pmatrix} \mathbf{\Sigma} & \mathbf{0}_{n \times n} \\ \mathbf{0}_{n \times n} & -\mathbf{\Sigma} \end{pmatrix} \frac{1}{\sqrt{2}} \begin{pmatrix} \mathbf{V}^T & \mathbf{U}^T \\ \mathbf{V}^T & -\mathbf{U}^T \end{pmatrix}.$$

Hence any symmetric eigen-solver can produce the SVD of a matrix \mathbf{A} without forming $\mathbf{A}\mathbf{A}^\top$ or $\mathbf{A}^\top\mathbf{A}$.

- Yet another relation to eigen-decomposition: If the eigendecomposition of a real symmetric matrix is $\mathbf{A} = \mathbf{W}\mathbf{\Lambda}\mathbf{W}^\top = \mathbf{W}\text{diag}(\lambda_1, \dots, \lambda_n)\mathbf{W}^\top$, then


$$\mathbf{A} = \mathbf{W}\mathbf{\Lambda}\mathbf{W}^\top = \mathbf{W} \begin{pmatrix} |\lambda_1| & & \\ & \ddots & \\ & & |\lambda_n| \end{pmatrix} \begin{pmatrix} \text{sgn}(\lambda_1) & & \\ & \ddots & \\ & & \text{sgn}(\lambda_n) \end{pmatrix} \mathbf{W}^\top$$

is the SVD of \mathbf{A} .

Applications of eigen-decomposition and SVD

- Principal components analysis (PCA). $\mathbf{X} \in \mathbb{R}^{n \times p}$ is a centered data matrix. Perform SVD $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ or equivalently $\mathbf{X}^\top\mathbf{X} = \mathbf{V}\mathbf{\Sigma}^2\mathbf{V}^\top$. The linear combinations $\tilde{\mathbf{x}}_i = \mathbf{X}\mathbf{v}_i$ are the principal components and have variance σ_i^2 . Usages:

1. Dimension reduction: reduce dimensionality p to $q \ll p$. Use top PCs $\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_q$ in downstream analysis.
2. Use PCs to adjust for confounding – a serious issue in association studies in large data sets.



- Low rank approximation, e.g, image/data compression.

Eckart-Young theorem:

$$\min_{\text{rank}(\mathbf{Y})=r} \|\mathbf{X} - \mathbf{Y}\|_F^2$$

is achieved by $\mathbf{Y} = \sum_{i=1}^r \sigma_i \mathbf{u}_i \mathbf{v}_i^T$ with optimal value $\sum_{i=1}^{r-1} \sigma_i^2$, where $(\sigma_i, \mathbf{u}_i, \mathbf{v}_i)$ are singular values and vectors of \mathbf{X} .

Gene Golub's 2691×598 picture requires $2691 \times 598 \times 6 = 9,655,308$ bytes (RGB 16 bit per channel). Rank 120 approximation requires $120 \times (2691 + 598) \times 6 = 2,368,080$ bytes. Rank 50 approximation requires $50 \times (2691 + 598) \times 6 = 986,700$ bytes. Rank 12 approximation requires $12 \times (2691 + 598) \times 8 = 236,808$ bytes.




Figure 2. Rank 12, 50, and 120 approximations to a rank 598 color photo of Gene Golub.

- Least squares, ridge regression, least squares over a sphere, ...
- Moore-Penrose (MP) inverse: Using thin SVD,

$$\mathbf{A}^+ = \mathbf{V} \mathbf{\Sigma}^+ \mathbf{U}^T,$$

where $\mathbf{\Sigma}^+ = \text{diag}(\sigma_1^{-1}, \dots, \sigma_r^{-1}, 0, \dots, 0)$, $r = \text{rank}(\mathbf{A})$. This is how the `ginv()` function is implemented in MASS package.

- Read KL and JM and do HW5 for some more applications.

One eigen-pair - power method

Assume $\mathbf{A} \in \mathbb{R}^{n \times n}$ symmetric.

- *Power method* iterates according to

$$\begin{aligned}\mathbf{x}^{(t)} &\leftarrow \frac{1}{\|\mathbf{A}\mathbf{x}^{(t-1)}\|_2} \mathbf{A}\mathbf{x}^{(t-1)} \\ \lambda_1^{(t)} &\leftarrow \mathbf{x}^{(t)\top} \mathbf{A}\mathbf{x}^{(t)}\end{aligned}$$

from an initial guess $\mathbf{x}^{(0)}$ of unit norm.

- Suppose we arrange $|\lambda_1| > |\lambda_2| \geq \dots \geq |\lambda_n|$ (the first inequality strict) with corresponding eigenvectors \mathbf{u}_i , and expand $\mathbf{x}^{(0)} = c_1\mathbf{u}_1 + \dots + c_n\mathbf{u}_n$, then

$$\begin{aligned}\mathbf{x}^{(t)} &= \frac{(\sum_i \lambda_i^t \mathbf{u}_i \mathbf{u}_i^\top) (\sum_i c_i \mathbf{u}_i)}{\|(\sum_i \lambda_i^t \mathbf{u}_i \mathbf{u}_i^\top) (\sum_i c_i \mathbf{u}_i)\|_2} \\ &= \frac{\sum_i c_i \lambda_i^t \mathbf{u}_i}{\|\sum_i c_i \lambda_i^t \mathbf{u}_i\|_2} \\ &= \frac{c_1 \mathbf{u}_1 + c_2 (\lambda_2/\lambda_1)^t \mathbf{u}_2 + \dots + c_n (\lambda_n/\lambda_1)^t \mathbf{u}_n}{\|c_1 \mathbf{u}_1 + c_2 (\lambda_2/\lambda_1)^t \mathbf{u}_2 + \dots + c_n (\lambda_n/\lambda_1)^t \mathbf{u}_n\|_2} \left(\frac{\lambda_1}{|\lambda_1|}\right)^t.\end{aligned}$$

Thus $\mathbf{x}^{(t)} - \frac{c_1 \mathbf{u}_1}{\|c_1 \mathbf{u}_1\|_2} \left(\frac{\lambda_1}{|\lambda_1|}\right)^t \rightarrow 0$ as $t \rightarrow \infty$. The convergence rate is $|\lambda_2|/|\lambda_1|$.

- $\mathbf{x}^{(t)\top} \mathbf{A}\mathbf{x}^{(t)}$ converges to λ_1 .
- *Inverse power method* for finding the eigenvalue of smallest absolute value: Substitute \mathbf{A} by \mathbf{A}^{-1} in the power method. (E.g., pre-compute LU or Cholesky of \mathbf{A}).
- *Shifted inverse power*: Substitute $(\mathbf{A} - \mu \mathbf{I})^{-1}$ in the power method. It converges to an eigenvalue close to the given μ .
- Initial guess of the desired eigenvalue can be obtain by Gerschgorin's circle theorem and so on.
- Power method also applies to asymmetric \mathbf{A} , e.g., PageRank problem costs $O(n)$ per iteration.

Top r eigen-pairs - orthogonal iteration

Generalization of power method to higher dimensional invariant subspace.

- *Orthogonal iteration*: Initialize $\mathbf{Q}^{(0)} \in \mathbb{R}^{n \times r}$ with orthonormal columns. For $t = 1, 2, \dots$,

$$\begin{aligned}\mathbf{Z}^{(t)} &\leftarrow \mathbf{A}\mathbf{Q}^{(t-1)} \quad (n^2r \text{ flops}) \\ \mathbf{Q}^{(t)}\mathbf{R}^{(t)} &\leftarrow \mathbf{Z}^{(t)} \quad (\text{QR factorization, } nr^2 - r^3/3 \text{ flops})\end{aligned}$$

- $\mathbf{Z}^{(t)}$ converges to the eigenspace of the largest r eigenvalues if they are real and separated from remaining spectrum. The convergence rate is $|\lambda_{r+1}|/|\lambda_r|$.

(Impractical) full eigen-decomposition - QR iteration

Assume $\mathbf{A} \in \mathbb{R}^{n \times n}$ symmetric.

- take $r = n$ in the orthogonal iteration. Then $\mathbf{Q}^{(t)}$ converges to the eigenspace \mathbf{U} of \mathbf{A} . This implies that

$$\mathbf{T}^{(t)} := \mathbf{Q}^{(t)T} \mathbf{A} \mathbf{Q}^{(t)}$$

converges to a diagonal form $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_n)$.

- Note how to compute $\mathbf{T}^{(t)}$ from $\mathbf{T}^{(t-1)}$

$$\begin{aligned}\mathbf{T}^{(t-1)} &= \mathbf{Q}^{(t-1)T} \mathbf{A} \mathbf{Q}^{(t-1)} = \mathbf{Q}^{(t-1)T} (\mathbf{A} \mathbf{Q}^{(t-1)}) = (\mathbf{Q}^{(t-1)T} \mathbf{Q}^{(t)}) \mathbf{R}^{(t)} \\ \mathbf{T}^{(t)} &= \mathbf{Q}^{(t)T} \mathbf{A} \mathbf{Q}^{(t)} = \mathbf{Q}^{(t)T} \mathbf{A} \mathbf{Q}^{(t-1)} \mathbf{Q}^{(t-1)T} \mathbf{Q}^{(t)} = \mathbf{R}^{(t)} (\mathbf{Q}^{(t-1)T} \mathbf{Q}^{(t)}).\end{aligned}$$

- *QR iteration*: Initialize $\mathbf{U}^{(0)} \in \mathbb{R}^{n \times n}$ orthogonal and set $\mathbf{T}^{(0)} = \mathbf{U}^{(0)T} \mathbf{A} \mathbf{U}^{(0)}$. For $t = 1, 2, \dots$


$$\begin{aligned}\mathbf{U}^{(t)}\mathbf{R}^{(t)} &\leftarrow \mathbf{T}^{(t-1)} \quad (\text{QR factorization}) \\ \mathbf{T}^{(t)} &\leftarrow \mathbf{R}^{(t)}\mathbf{U}^{(t)}\end{aligned}$$

- QR iteration is expensive: $O(n^3)$ per iteration and linear convergence rate.

QR algorithm for symmetric eigen-decomposition

Assume $\mathbf{A} \in \mathbb{R}^{n \times n}$ symmetric.

- As done in LAPACK: `eigen()` in R, `eig()` in Matlab
- Idea: Tri-diagonalization (by Householder) + QR iteration on the tri-diagonal system with implicit shift
 - Step 1: Householder tri-diagonalization: $4n^3/3$ for eigenvalues only, $8n^3/3$ for both eigenvalues and eigenvectors. (Why can't we apply Householder to make it diagonal directly?)



- Step 2: QR iteration on the tridiagonal matrix. Implicit shift accelerates convergence rate. On average 1.3-1.6 QR iteration per eigenvalue, $\sim 20n$ flops per QR iteration. So total operation count is about $30n^2$. Eigenvectors need an extra of about $6n^3$ flops.

	Eigenvalue	Eigenvector
Householder reduction	$4n^3/3$	$4n^3/3$
QR with implicit shift	$\sim 30n^2$	$\sim 6n^3$

- Don't request eigenvectors unless necessary: set `only.values = TRUE` when calling `eigen()` in R.
- Remark: there are at least two alternative ways (other than QR with implicit shift) to solve the symmetric tridiagonal eigenproblem

- Bisection: good when only selected portions of eigensystem are needed
 - Divide and conquer: good for parallel implementation
- The *unsymmetric QR algorithm* obtains the real Schur decomposition of an asymmetric matrix \mathbf{A} .

14 Lecture 14, Oct 16

Announcements

- HW4 returned. Feedback:
 - Key to the success of iterative methods is fast matrix-vector multiplication. Here the structure in \mathbf{P}^T or $\mathbf{I} - \mathbf{P}^T$ is “sparse + low rank”. Say in power method, if we first form the matrix \mathbf{P}^T and then iterate according to $\mathbf{x}^{(t)} \leftarrow \mathbf{P}^T \mathbf{x}^{(t-1)}$, it costs n^2 flops per iteration. If we keep the “sparse + low rank” structure and iterate according to

$$\mathbf{x}^{(t)} \leftarrow \mathbf{A}^T(\text{diag}(\mathbf{r}^+) \mathbf{x}^{(t-1)}) + \mathbf{1}_n(\mathbf{z}^T \mathbf{x}^{(t-1)}),$$

it costs $3n + \text{nnz}(\mathbf{A}) \ll n^2$ flops, where $\text{nnz}(\mathbf{A})$ is the number of non-zero elements in \mathbf{A} .

- Sketch of solution: <http://hua-zhou.github.io/teaching/st758-2014fall/hw04sol.html>
- Some run times in MATLAB: using vs ignoring “sparse + low rank” structure.

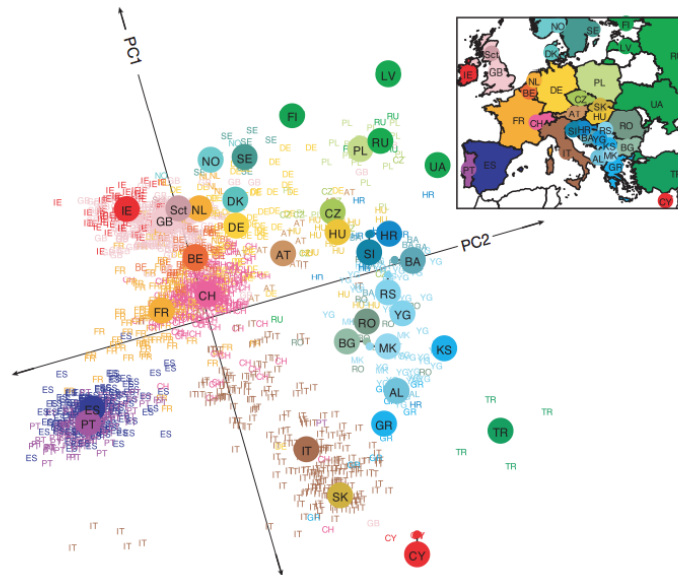
Matrix	Size	Sparsity	svt (fh input)	svds
bfwb398	398	0.9816	0.0176(0.0011)	0.0408(0.0009)
rdb800l	800	0.9928	0.0240(0.0005)	0.2115(0.0014)
tols1090	1090	0.9970	0.0780(0.0009)	0.9396(0.0079)
mhd4800b	4800	0.9988	0.0471(0.0001)	5.6700(0.0166)
cryg10000	10000	0.9995	0.1909(0.0022)	44.2213(0.4373)

Table 3: Top 6 singular values and vectors of “sparse + low rank” matrices by **svt** and **svds**. Structured matrices are formed by adding a random rank-10 matrix to the original sparse test matrix. Reported are the average run time (in seconds) and standard error (in parentheses) based on 10 simulation replicates.

- HW5 due next Tue Oct 21.

Last time

- Linear algebra review: SVD
- Applications of eigen-decomposition and SVD:
PCA, data/image compression, least squares and its variants, MP inverse, stationary distribution of Markov chains (Google PageRank), ...
- Answer to Caleb's question: Following picture is from the article "Genes mirror geography within Europe" by Novembre et al. (2008) published in *Nature* <http://www.nature.com/nature/journal/v456/n7218/full/nature07331.html>.




Use of PCA to adjust for confounding in modern genetic studies is proposed in the paper "Principal components analysis corrects for stratification in genome-wide association studies" by Price et al. (2006) published in *Nature Genetics* <http://www.nature.com/ng/journal/v38/n8/full/ng1847.html>. It has been cited 3346 times as of Oct 14, 2014.

- Algorithms for eigen-decomposition:
power method (top eigen-pair), orthogonal iteration (top r eigen-pairs), QR iteration for full eigen-decomposition ($O(n^4)$ flops), Householder tri-diagonalization + QR iteration with implicit shift ($O(n^3)$ flops).

Today

- SVD algorithm

Algorithm for singular value decomposition (SVD)



Gene Golub's license plate, photographed by Professor P. M. Kooenbeig of Leiden University

Assume $\mathbf{A} \in \mathbb{R}^{m \times n}$ and we seek the SVD $\mathbf{A} = \mathbf{U}\mathbf{D}\mathbf{V}^T$.

- “Golub-Kahan-Reinsch” algorithm:
 - Stage 1: Transform \mathbf{A} to an upper bidiagonal form \mathbf{B} (by Householder).

$$U_B^T A V_B = \begin{bmatrix} B \\ 0 \end{bmatrix} \quad B = \begin{bmatrix} d_1 & f_1 & & \cdots & 0 \\ 0 & d_2 & & & \vdots \\ & & \ddots & & \\ & & & \ddots & \\ \vdots & & & & \ddots & f_{n-1} \\ 0 & \cdots & & 0 & & d_n \end{bmatrix} \in \mathbb{R}^{n \times n}$$

$$\begin{array}{ccc}
\begin{bmatrix} \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \end{bmatrix} & \xrightarrow{U_1} & \begin{bmatrix} \times & \times & \times & \times \\ 0 & \times & \times & \times \\ 0 & \times & \times & \times \\ 0 & \times & \times & \times \\ 0 & \times & \times & \times \end{bmatrix} & \xrightarrow{V_1} \\
\\
\begin{bmatrix} \times & \times & 0 & 0 \\ 0 & \times & \times & \times \\ 0 & \times & \times & \times \\ 0 & \times & \times & \times \\ 0 & \times & \times & \times \end{bmatrix} & \xrightarrow{U_2} & \begin{bmatrix} \times & \times & 0 & 0 \\ 0 & \times & \times & \times \\ 0 & 0 & \times & \times \\ 0 & 0 & \times & \times \\ 0 & 0 & \times & \times \end{bmatrix} & \xrightarrow{V_2} \\
\\
\begin{bmatrix} \times & \times & 0 & 0 \\ 0 & \times & \times & 0 \\ 0 & 0 & \times & \times \\ 0 & 0 & \times & \times \\ 0 & 0 & \times & \times \end{bmatrix} & \xrightarrow{U_3} & \begin{bmatrix} \times & \times & 0 & 0 \\ 0 & \times & \times & 0 \\ 0 & 0 & \times & \times \\ 0 & 0 & 0 & \times \\ 0 & 0 & 0 & \times \end{bmatrix} & \xrightarrow{U_4} & \begin{bmatrix} \times & \times & 0 & 0 \\ 0 & \times & \times & 0 \\ 0 & 0 & \times & \times \\ 0 & 0 & 0 & \times \\ 0 & 0 & 0 & 0 \end{bmatrix}
\end{array}$$

– Stage 2: Apply implicit-shift QR step to the tridiagonal matrix $\mathbf{B}^T \mathbf{B}$ *implicitly*.

- See Golub and Van Loan (1996, Section 8.6) for more details.
- $4m^2n + 8mn^2 + 9n^3$ flops for a tall ($m \geq n$) matrix.
- `svd()` in R and Matlab: wrapper of the `_GESVD` subroutine in LAPACK.

15 Lecture 15, Oct 21

Announcements

- HW5 due today: submit hardcopy + email code
- Homework due date change to Thu?
- Spaces around “=” in function call?

Last time

- HW4 (PageRank algorithm). Ilse Ipsen’s slides
http://www4.ncsu.edu/~ipsen/ps/slides_imacs.pdf
- SVD (Golub-Kahan-Riensch) algorithm

Today

- Iterative methods (Lanszos and Arnoldi methods) for huge, structured \mathbf{A}
- Jacobi method for eigen-decomposition (parallel computing)
- Generalized eigen-problem: $\mathbf{Ax} = \lambda\mathbf{Bx}$ (PLS, SIR, CCA, ...)
- Variants of least squares problems (self-study)
- Concluding remarks for numerical linear algebra

Lanczos/Arnoldi iterative method for top eigen-pairs

- Motivation
 - Consider the Google PageRank problem. We want to find the top left eigenvector of the transition matrix

$$\mathbf{P} = p\mathbf{R}^+\mathbf{A} + z\mathbf{1}_n^\top,$$

where $\mathbf{R} = \text{diag}(r_1, \dots, r_n)$ and $z_j = (1 - p)/n$ if $r_i > 0$ and $1/n$ if $r_i = 0$. Suppose there are $n \approx 1$ billion web pages.


The (unsymmetric) QR algorithm will take order

$$\frac{(1 \times 10^9)^3}{10^{12}} \approx 3.33 \times 10^{14} \text{ seconds} \approx 1 \times 10^7 \text{ years}$$

on a tera-flop supercomputer!

- Consider adjusting for confounding by PCA in modern GWAS (genome-wide association studies). We want to find the top singular values/vectors of a genotype matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$, where $n \sim 10^3$ and $p \sim 10^6$.
- Krylov subspace methods are the state-of-art iterative method for obtaining the top eigen-values/vectors or singular values/vectors of large sparse or structured matrices.
- Lanczos method: top eigen-pairs of a large *symmetric* matrix.
- Arnoldi method: top eigen-pairs of a large *asymmetric* matrix.
- Both methods are also adapted to obtain top singular values/vectors of large sparse or structured matrices.
- We will give a brief overview of these methods together with the conjugate gradient method for solving large linear system.
- `eigs()` and `svds()` in `Matlab` are wrappers of the ARPACK package. No native functions in R (?).

Jacobi method for symmetric eigen-decomposition (KL 8.2)



Assume $\mathbf{A} \in \mathbb{R}^{n \times n}$ is symmetric and we seek the eigen-decomposition $\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$.

- Idea: Systematically reduce off-diagonal entries

$$\text{off}(\mathbf{A}) = \sum_i \sum_{j \neq i} a_{ij}^2$$

by Jacobi rotations.

- Jacobi/Givens rotations:

$$\mathbf{J}(p, q, \theta) = \begin{pmatrix} 1 & & 0 & & 0 \\ \vdots & \ddots & \vdots & & \vdots \\ 0 & & \cos(\theta) & & \sin(\theta) \\ \vdots & & \vdots & \ddots & \vdots \\ 0 & & -\sin(\theta) & & \cos(\theta) \\ \vdots & & \vdots & & \vdots \\ 0 & & 0 & & 0 & \ddots & \vdots \\ & & & & & & & 1 \end{pmatrix},$$

$\mathbf{J}(p, q, \theta)$ is orthogonal.

- Consider $\mathbf{B} = \mathbf{J}^T \mathbf{A} \mathbf{J}$. \mathbf{B} preserves the symmetry and eigenvalues of \mathbf{A} .

Taking

$$\begin{cases} \tan(2\theta) = 2a_{pq}/(a_{qq} - a_{pp}) & \text{if } a_{pp} \neq a_{qq} \\ \theta = \pi/4 & \text{if } a_{pp} = a_{qq} \end{cases}$$

forces $b_{pq} = 0$.

- Since orthogonal transform preserves Frobenius norm, we have

$$b_{pp}^2 + b_{qq}^2 = a_{pp}^2 + a_{qq}^2 + 2a_{pq}^2.$$

(Just check the 2-by-2 block)

- Since $\|\mathbf{A}\|_F = \|\mathbf{B}\|_F$, this implies that the off-diagonal part

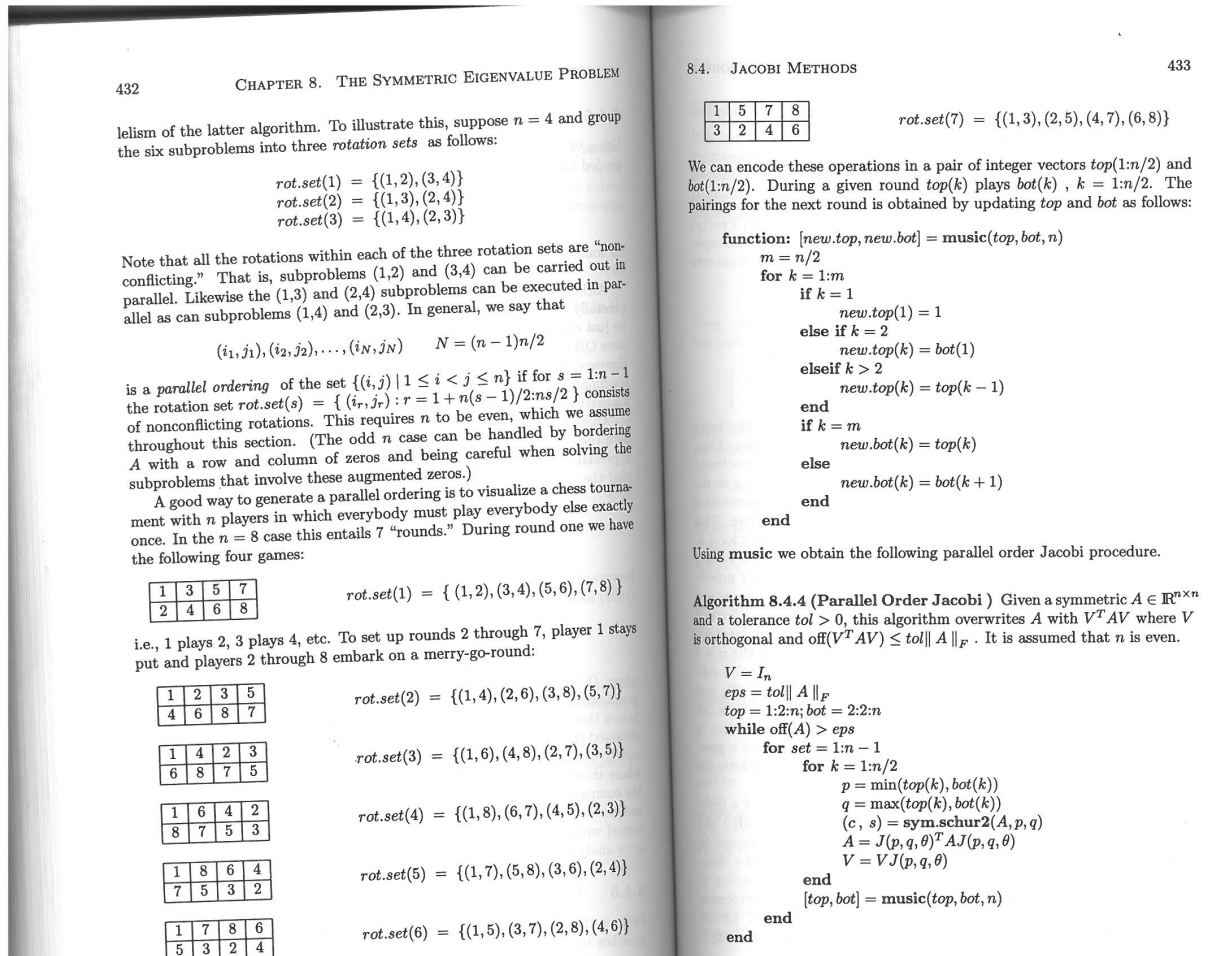
$$\text{off}(\mathbf{B}) = \text{off}(\mathbf{A}) - 2a_{pq}^2$$

is decreased whenever $a_{pq} \neq 0$.

- One Jacobi rotation costs $O(n)$ flops.
- *Classical Jacobi*: search for the largest $|a_{ij}|$ at each iteration.
- $\text{off}(\mathbf{A}) \leq n(n-1)a_{ij}^2$ and $\text{off}(\mathbf{B}) = \text{off}(\mathbf{A}) - 2a_{ij}^2$ together implies

$$\text{off}(\mathbf{B}) \leq \left(1 - \frac{2}{n(n-1)}\right) \text{off}(\mathbf{A}).$$

- In practice, cyclic-by-row implementation, to avoid the costly $O(n^2)$ search in the classical Jacobi.
- Jacobi method attracts a lot recent attention because of its rich inherent parallelism.
- *Parallel Jacobi*: “merry-go-round” to generate parallel ordering.



Generalized eigen-problem

- Generalized eigen-problem: $\mathbf{A}\mathbf{x} = \lambda\mathbf{B}\mathbf{x}$, where \mathbf{A} psd and \mathbf{B} pd.
- Applications: partial least squares (PLS), sliced inverse regression (SIR), canonical correlation analysis (CCA).
- Method 1: $\mathbf{B}^{-1}\mathbf{A}\mathbf{x} = \lambda\mathbf{x}$. Non-symmetric eigen-problem \ominus .
- Method 2: Cholesky $\mathbf{B} = \mathbf{L}\mathbf{L}^\top$. Then $\mathbf{L}^{-1}\mathbf{A}\mathbf{L}^{-\top}\mathbf{y} = \lambda\mathbf{y}$ where $\mathbf{y} = \mathbf{L}^\top\mathbf{x}$.
- Method 3 (most numerically stable, \mathbf{B} can be rank deficient): QZ algorithm.
- `eig()` and `qz()` in `Matlab` implement QZ. No native function in R?

Generalized singular value decomposition

- $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{B} \in \mathbb{R}^{p \times n}$. Then there exists orthogonal $\mathbf{U} \in \mathbb{R}^{m \times m}$ and $\mathbf{V} \in \mathbb{R}^{p \times p}$ and an invertible $\mathbf{X} \in \mathbb{R}^{n \times n}$ such that

$$\begin{aligned} \mathbf{U}^\top \mathbf{A} \mathbf{X} &= \mathbf{C} = \text{diag}(c_1, \dots, c_n), & c_i &\geq 0 \\ \mathbf{V}^\top \mathbf{B} \mathbf{X} &= \mathbf{S} = \text{diag}(s_1, \dots, s_q), & s_i &\geq 0, \end{aligned}$$

where $q = \min\{p, n\}$.

- Applications: quadratically inequality-constrained least squares problem (LSQI).
- `gsvd()` in `Matlab` implements generalized SVD. No native function in R?

In the zoo of least squares (self-study)

Weighted least squares

- In weighted least squares, we minimize $\sum_{i=1}^n w_i (y_i - \mathbf{x}_i^\top \boldsymbol{\beta})^2$, where $w_i > 0$ are observation weights.
- Let $\mathbf{W} = \text{diag}(w_1, \dots, w_n)$. Then the criterion is $\|\mathbf{W}^{1/2}\mathbf{y} - \mathbf{W}^{1/2}\mathbf{X}\boldsymbol{\beta}\|_2^2$, which can be solved by standard methods for least squares with $\tilde{\mathbf{y}} = \mathbf{W}^{1/2}\mathbf{y}$ and $\tilde{\mathbf{X}} = \mathbf{W}^{1/2}\mathbf{X}$.

General least squares

- In Aitken model: $E(\mathbf{y}) = \mathbf{X}\boldsymbol{\beta}$, $\text{Cov}(\mathbf{y}) = \sigma^2\mathbf{V}$, where \mathbf{V} is a positive semidefinite matrix. We minimize the generalized least squares criterion

$$(\mathbf{y} - \mathbf{X}\boldsymbol{\beta})^T \mathbf{M}^{-1} (\mathbf{y} - \mathbf{X}\boldsymbol{\beta}),$$

where $\mathbf{M} \in \mathbb{R}^{n \times n}$ is some positive semidefinite matrix, e.g., $\mathbf{M} = \mathbf{V}$ for non-singular \mathbf{V} or $\mathbf{M} = \mathbf{V} + \mathbf{X}\mathbf{X}^T$ for singular \mathbf{V} .

- Let $\mathbf{M} = \mathbf{B}\mathbf{B}^T$ for some $\mathbf{B} \in \mathbb{R}^{n \times n}$ (e.g., the Cholesky factor). One approach is to minimize

$$\|\mathbf{B}^{-1}(\mathbf{y} - \mathbf{X}\boldsymbol{\beta})\|_2^2.$$

Unfortunately, when \mathbf{B} is poorly conditioned (or even not invertible), the procedure produces a poor solution.

- Paige's method. The generalized least squares problem is equivalent to

$$\begin{aligned} & \text{minimize} \quad \mathbf{v}^T \mathbf{v} \\ & \text{subject to} \quad \mathbf{X}\boldsymbol{\beta} + \mathbf{B}\mathbf{v} = \mathbf{y}. \end{aligned}$$

To solve this problem, first compute the QR of \mathbf{X}

$$\mathbf{X} = (\mathbf{Q}_1, \mathbf{Q}_2) \begin{pmatrix} \mathbf{R}_1 \\ \mathbf{0} \end{pmatrix}.$$

Compute another QR for the (flat) matrix $\mathbf{Q}_2^T \mathbf{B}$ such that

$$\mathbf{Q}_2^T \mathbf{B} = (\mathbf{0}, \mathbf{S}) \begin{pmatrix} \mathbf{Z}_1^T \\ \mathbf{Z}_2^T \end{pmatrix},$$

where \mathbf{S} is upper triangular and $(\mathbf{Z}_1, \mathbf{Z}_2) \in \mathbb{R}^{n \times n}$ is orthogonal. Then the constraint becomes

$$\begin{pmatrix} \mathbf{R}_1 \\ \mathbf{0} \end{pmatrix} \boldsymbol{\beta} + \begin{pmatrix} \mathbf{Q}_1^T \mathbf{B} \mathbf{Z}_1 & \mathbf{Q}_1^T \mathbf{B} \mathbf{Z}_2 \\ \mathbf{0} & \mathbf{S} \end{pmatrix} \begin{pmatrix} \mathbf{Z}_1^T \mathbf{v} \\ \mathbf{Z}_2^T \mathbf{v} \end{pmatrix} = \begin{pmatrix} \mathbf{Q}_1^T \mathbf{y} \\ \mathbf{Q}_2^T \mathbf{y} \end{pmatrix}.$$

From the bottom half we can solve for \mathbf{v} from the equation (how?)

$$\mathbf{S} \mathbf{Z}_2^T \mathbf{v} = \mathbf{Q}_2^T \mathbf{y}.$$

Then we solve for $\boldsymbol{\beta}$ from the equation

$$\mathbf{R}_1 \boldsymbol{\beta} = \mathbf{Q}_1^T \mathbf{y} - (\mathbf{Q}_1^T \mathbf{B} \mathbf{Z}_1 \mathbf{Z}_1^T + \mathbf{Q}_1^T \mathbf{B} \mathbf{Z}_2 \mathbf{Z}_2^T) \mathbf{v} = \mathbf{Q}_1^T \mathbf{y} - \mathbf{Q}_1^T \mathbf{B} \mathbf{Z}_2 (\mathbf{Z}_2^T \mathbf{v}).$$

- Paige's method also works for singular \mathbf{X} and \mathbf{B} (using QR with column pivoting).
- MATLAB's `lscov()` function implements Paige's method for singular covariance \mathbf{V} . No R implementation (?)

Ridge regression

- In ridge regression, we minimize

$$\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda\|\boldsymbol{\beta}\|_2^2,$$

where λ is a tuning parameter.

- Ridge regression by augmented linear regression. Ridge regression problem is equivalent to

$$\left\| \begin{pmatrix} \mathbf{y} \\ \mathbf{0}_p \end{pmatrix} - \begin{pmatrix} \mathbf{X} \\ \sqrt{\lambda}\mathbf{I}_p \end{pmatrix} \boldsymbol{\beta} \right\|_2^2.$$

Therefore any methods for linear regression can be applied.

- Ridge regression by method of normal equation. The normal equation for the ridge problem is

$$(\mathbf{X}^T \mathbf{X} + \lambda \mathbf{I}_p) \boldsymbol{\beta} = \mathbf{X}^T \mathbf{y}.$$

Therefore Cholesky or sweep can be used.

- Ridge regression by SVD. If we obtain the (thin) SVD of \mathbf{X}

$$\mathbf{X} = \mathbf{U} \boldsymbol{\Sigma}_{p \times p} \mathbf{V}^T.$$

Then the normal equation reads

$$(\boldsymbol{\Sigma}^2 + \lambda \mathbf{I}_p) \mathbf{V}^T \boldsymbol{\beta} = \boldsymbol{\Sigma} \mathbf{U}^T \mathbf{y}$$

and we get

$$\hat{\boldsymbol{\beta}}(\lambda) = \sum_{i=1}^p \frac{\sigma_i \mathbf{u}_i^T \mathbf{y}}{\sigma_i^2 + \lambda} \mathbf{v}_i = \sum_{i=1}^r \frac{\sigma_i \mathbf{u}_i^T \mathbf{y}}{\sigma_i^2 + \lambda} \mathbf{v}_i, \quad r = \text{rank}(\mathbf{X}).$$

It is clear that

$$\lim_{\lambda \rightarrow 0} \widehat{\boldsymbol{\beta}}(\lambda) = \widehat{\boldsymbol{\beta}}_{\text{OLS}}$$

and $\|\widehat{\boldsymbol{\beta}}(\lambda)\|_2$ is monotone decreasing as λ increases.

- Only one SVD is needed for all λ (!), in contrast to the method of augmented linear regression, Cholesky, or sweep.

Least squares over a sphere

- Ridge regression “shrinks” the solution via penalty. Alternatively we can simply fit a least squares problem subject to the constraint that the solution lives in a sphere

$$\begin{aligned} &\text{minimize} && \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 \\ &\text{subject to} && \|\boldsymbol{\beta}\|_2 \leq \alpha. \end{aligned}$$

- Suppose we obtain the (thin) SVD $\mathbf{X} = \mathbf{U}\boldsymbol{\Sigma}_{p \times p}\mathbf{V}^T$. If the ordinary least squares solution

$$\widehat{\boldsymbol{\beta}}_{\text{OLS}} = \sum_{i=1}^r \frac{\mathbf{u}_i^T \mathbf{y}}{\sigma_i} \mathbf{v}_i$$

has ℓ_2 norm less than α , then we are done. If not, we use the method of Lagrangian multipliers

$$\psi(\boldsymbol{\beta}, \lambda) = \frac{1}{2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \frac{\lambda}{2} (\|\boldsymbol{\beta}\|_2^2 - \alpha^2).$$

Setting the gradient to 0, we have the shifted normal equation

$$(\mathbf{X}^T \mathbf{X} + \lambda \mathbf{I}_p) \boldsymbol{\beta} = \mathbf{X}^T \mathbf{y},$$

which has solution

$$\widehat{\boldsymbol{\beta}}(\lambda) = \sum_{i=1}^r \frac{\sigma_i \mathbf{u}_i^T \mathbf{y}}{\sigma_i^2 + \lambda} \mathbf{v}_i.$$

We need to choose the λ such that $\|\widehat{\boldsymbol{\beta}}(\lambda)\|_2 = \alpha$. That is we need to find the (unique) zero of the function

$$f(\lambda) = \|\widehat{\boldsymbol{\beta}}(\lambda)\|_2^2 - \alpha^2 = \sum_{i=1}^r \left(\frac{\sigma_i \mathbf{u}_i^T \mathbf{y}}{\sigma_i^2 + \lambda} \right)^2 - \alpha^2.$$

This is easily achieved by Newton’s or other methods.

Least squares with equality constraints

- In many applications, there are *a priori* constraints on the regression parameters. Let's consider how to solve linear regression with equality constraints (LSE)

$$\begin{aligned} & \text{minimize} && \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 \\ & \text{subject to} && \mathbf{B}\boldsymbol{\beta} = \mathbf{d}. \end{aligned}$$

- LSE by QR. First compute QR of $\mathbf{B}^T \in \mathbb{R}^{p \times m}$

$$\mathbf{B}^T = \mathbf{Q} \begin{pmatrix} \mathbf{R} \\ \mathbf{0} \end{pmatrix}$$

and set

$$\mathbf{X}\mathbf{Q} = (\mathbf{X}_1, \mathbf{X}_2) \quad \text{and} \quad \mathbf{Q}^T\boldsymbol{\beta} = \begin{pmatrix} \boldsymbol{\beta}_1 \\ \boldsymbol{\beta}_2 \end{pmatrix}.$$

Then the original minimization problem becomes

$$\begin{aligned} & \text{minimize} && \|\mathbf{y} - \mathbf{X}_1\boldsymbol{\beta}_1 - \mathbf{X}_2\boldsymbol{\beta}_2\|_2^2 \\ & \text{subject to} && \mathbf{R}^T\boldsymbol{\beta}_1 = \mathbf{d}. \end{aligned}$$

Now $\boldsymbol{\beta}_1$ is determined from the constraint $\mathbf{R}^T\boldsymbol{\beta}_1 = \mathbf{d}$ and $\boldsymbol{\beta}_2$ is solved from the unconstrained least squares problem

$$\text{minimize} \quad \|(\mathbf{y} - \mathbf{X}_1\boldsymbol{\beta}_1) - \mathbf{X}_2\boldsymbol{\beta}_2\|_2^2.$$

Finally we recover the solution from

$$\boldsymbol{\beta} = \mathbf{Q} \begin{pmatrix} \boldsymbol{\beta}_1 \\ \boldsymbol{\beta}_2 \end{pmatrix}.$$

- LSE by augmented system. Define the Lagrangian function

$$\phi(\boldsymbol{\beta}, \boldsymbol{\lambda}) = \frac{1}{2}\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 - \boldsymbol{\lambda}^T(\mathbf{B}\boldsymbol{\beta} - \mathbf{d}).$$

Setting gradient to zero yields

$$\begin{aligned} \mathbf{X}^T\mathbf{X}\boldsymbol{\beta} - \mathbf{B}^T\boldsymbol{\lambda} &= \mathbf{X}^T\mathbf{y} \\ \mathbf{B}\boldsymbol{\beta} &= \mathbf{d}, \end{aligned}$$


suggesting the augmented system

$$\begin{pmatrix} \mathbf{X}^T \mathbf{X} & \mathbf{B}^T \\ \mathbf{B} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \boldsymbol{\beta} \\ -\lambda \end{pmatrix} = \begin{pmatrix} \mathbf{X}^T \mathbf{y} \\ \mathbf{d} \end{pmatrix}.$$

This linear system is non-singular when \mathbf{X} and \mathbf{B} have full rank and can be solved by Cholesky, sweep, and so on.

- LSE by generalized SVD.

Total least squares (TLS)



TLS considers the case both predictors and observations are subject to errors. It is solved by SVD. Read KL 9.3.6 if interested.

Tikhonov regularization

Tikhonov regularization is an extension of the ridge regression

$$\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda\|\mathbf{B}\boldsymbol{\beta}\|_2^2,$$

where $\mathbf{B} \in \mathbb{R}^{m \times p}$ is a fixed regularization matrix and λ is a tuning parameter. It is solved by the generalized singular value decomposition (GSVD).

Least squares with quadratic inequality constraint (LSQI)

Least squares with quadratic inequality constraint (LSQI) minimizes the least squares criterion over a hyper-ellipsoid:

$$\begin{aligned} &\text{minimize} && \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 \\ &\text{subject to} && \|\mathbf{B}\boldsymbol{\beta}\|_2 \leq \alpha, \end{aligned}$$

where $\mathbf{B} \in \mathbb{R}^{m \times p}$ is a fixed regularization matrix. It is solved by the generalized singular value decomposition (GSVD). See Golub and Van Loan (1996, Section 2.1.1).

Concluding remarks on numerical linear algebra

- Numerical linear algebra forms the building blocks of most computation we do. Most lines of our code are numerical linear algebra.
- Be flop and memory aware.

The form of a mathematical expression and the way the expression should be evaluated in actual practice may be quite different.





- Be alert to problem structure and make educated choice of software/algorithm.

The structure should be exploited whenever solving a problem.



- Do not write your own matrix computation routines unless for good reason. Utilize BLAS and LAPACK as much as possible!
- In contrast, for optimization, often we need to devise problem specific optimization routines, or even “mix and match” them.

Some useful reference books on (numerical) linear algebra



- Linear algebra books: Magnus and Neudecker (1999), Horn and Johnson (1985), Harville (1997), Gentle (2007)




- Golub and Van Loan (1996): “Bible” in numerical linear algebra.



- Lawson and Hanson (1987) and Björck (1996): classical monographs on solving least squares problems.



- Saad (2003): standard reference for iterative methods



16 Lecture 16, Oct 23

Announcements

- HW5 returned. Feedback:
 - What's the likelihood of two independently written code looking like this?

```
data <- read.table(file = "http://hua-zhou.github.io/teaching/st758-2014fall/hw05sol.dat")
y <- data[, 1]
n <- length(y)
x <- data[, -1]
x <- as.matrix(x)
l <- rep(1, n)
X <- cbind(l, x)
s <- svd(X)
d <- s$d
U <- s$u
V <- s$v
D <- d^2
U_y <- t(U) %*% y
z <- matrix(0, nrow(U_y), 1)
z <- d * U_y
ridge_soln <- function(k)
{Dmodified <- D + k
Dmodified <- 1 / (Dmodified) ^ 2
z <- Dmodified * z ^ 2
return(sqrt(sum(z)))
}
lambda <- seq(5, 100, by=5)
y1 <- sapply(1:20, function(x) ridge_soln(5 * x))
plot(lambda, y1, col = "red", main = "RIDGE SOLUTION", xlab = expr = "Values")

data <- read.table("longley.txt")
x <- data[, -1]
y <- data[, 1]
n <- length(y)
x <- as.matrix(x)
l <- rep(1, n)
X <- cbind(l, x) # adding intercept
s <- svd(X)
D <- s$d # array of diagonal elements
D_matrix <- diag(D)
no_of_elmts <- ncol(t(D))
U <- s$u
V <- s$v
D2 <- D ^ 2
Uty <- t(U) %*% y
z <- matrix(0, nrow(Uty), 1)
z <- D * Uty

ridgeBeta <- function(k)
{D2new <- D2 + k
D2new <- 1 / (D2new) ^ 2
z <- D2new * z ^ 2
return(sqrt(sum(z)))
}
lambda <- seq(5, 100, by = 5)
w <- sapply(1:20, function(x) ridgeBeta(5 * x))
plot(lambda, w, col = "blue", main = "Ridge Betahat norm",
ylab = "Values")
```

Copying code from any source without acknowledgement is plagiarism.

First time, hw score is 0; second time, course grade is F.

- Q2 (ridge regression): avoid loops. Vectorize code (matrix/vector operation) > `apply`, `sapply` > for loop.
- Organize your code into functions.
- Sketch of solution: <http://hua-zhou.github.io/teaching/st758-2014fall/hw05sol.html>

- HW6 posted and due Nov 6. Start early.
- ST790: Advanced Statistical Computing.

Last time

- Iterative methods (Lanszos and Arnoldi methods) for huge, structured \mathbf{A}

- Jacobi method for eigen-decomposition (parallel computing)
- Generalized eigen-problem: $\mathbf{Ax} = \lambda\mathbf{Bx}$ (PLS, SIR, CCA, ...) and generalized singular value decomposition
- Variants of least squares problem
- Concluding remarks for numerical linear algebra

Today

- MLE
- General optimization theory

MLE (as a motivation for optimization)

A great idea due to Fisher in 20s, and made rigorous by Cramer and others in 40s.

- Notations:
 - Density: $f(\mathbf{x}|\boldsymbol{\theta})$, where $\boldsymbol{\theta} \in \Theta \subset \mathbb{R}^p$
 - Log-likelihood function: $L(\boldsymbol{\theta}) = \ln f(\mathbf{x}|\boldsymbol{\theta})$
 - (Column) Gradient/score vector: $\nabla L(\boldsymbol{\theta}) \in \mathbb{R}^{p \times 1}$
 - Differential: $dL(\boldsymbol{\theta}) = [\nabla L(\boldsymbol{\theta})]^\top \in \mathbb{R}^{1 \times p}$
 - Hessian: $d^2L(\boldsymbol{\theta}) = \nabla^2 L(\boldsymbol{\theta})$
 - Observed information matrix: $-d^2L(\boldsymbol{\theta})$
 - Expected (Fisher) information matrix: $\mathbf{I}(\boldsymbol{\theta}) = \mathbf{E}_{\boldsymbol{\theta}}[-d^2L(\boldsymbol{\theta})]$
 - Given iid observations $\mathbf{x}_1, \dots, \mathbf{x}_n$ from $f(\cdot|\boldsymbol{\theta})$,

$$L_n(\boldsymbol{\theta}) = \sum_{i=1}^n \ln f(\mathbf{x}_i|\boldsymbol{\theta})$$

- Maximum likelihood estimator (MLE):

$$\hat{\boldsymbol{\theta}}_{\text{MLE}} = \operatorname{argmax}_{\boldsymbol{\theta}} L_n(\boldsymbol{\theta})$$

- Consistency of MLE

- Under the true parameter value $\boldsymbol{\theta}_0$,

$$M_n(\boldsymbol{\theta}) = \frac{1}{n}[L_n(\boldsymbol{\theta}) - L_n(\boldsymbol{\theta}_0)]$$

$$\rightarrow M(\boldsymbol{\theta}) = \mathbf{E}_{\boldsymbol{\theta}_0}[\ln f(\mathbf{X}|\boldsymbol{\theta}) - \ln f(\mathbf{X}|\boldsymbol{\theta}_0)]$$

for all $\boldsymbol{\theta}$ almost surely.

- Note that $M(\boldsymbol{\theta})$ is the negative Kullback-Leibler divergence between distribution at $\boldsymbol{\theta}$ and distribution at $\boldsymbol{\theta}_0$.

Assuming *identifiability*, by the *information inequality*, $M(\boldsymbol{\theta})$ achieves maximum uniquely at $\boldsymbol{\theta}_0$. We hope the MLE

$$\hat{\boldsymbol{\theta}}_n = \operatorname{argmax}_{\boldsymbol{\theta}} M_n(\boldsymbol{\theta})$$

converges to

$$\boldsymbol{\theta}_0 = \operatorname{argmax}_{\boldsymbol{\theta}} M(\boldsymbol{\theta}).$$

- Need *uniform convergence* of $M_n(\boldsymbol{\theta})$ to $M(\boldsymbol{\theta})$, i.e.,

$$\sup_{\Theta} |M_n(\boldsymbol{\theta}) - M(\boldsymbol{\theta})|$$

converges to 0 in probability. A set of sufficient conditions for uniform convergence:

- * compactness of the parameter space Θ
 - * continuity of $M(\boldsymbol{\theta})$ in $\boldsymbol{\theta}$ for any \boldsymbol{x}
 - * $M(\boldsymbol{\theta})$ dominated by an integrable function
- Example of non-uniform convergence: $f_n(x) = 1_{\{n, n+1\}}$ (or a triangle on $[n, n+1]$ if we want f_n to be continuous). $f_n \rightarrow f \equiv 0$ pointwise but not uniformly.

- Asymptotic normality of MLE

- Assume $\hat{\boldsymbol{\theta}}_n$ is consistent for $\boldsymbol{\theta}_0$.
- Taylor expansion on $\mathbf{0}_p = \frac{1}{n} \nabla L_n(\hat{\boldsymbol{\theta}}_n)$ gives

$$\mathbf{0}_p = \frac{1}{n} \nabla L_n(\boldsymbol{\theta}_0) + \left[\frac{1}{n} d^2 L_n(\boldsymbol{\theta}_0) \right] (\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0)$$

$$+ \frac{1}{2} [\mathbf{I}_p \otimes (\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0)^\top] \left[\frac{1}{n} D d^2 L_n(\tilde{\boldsymbol{\theta}}_n) \right] (\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0),$$

where $\tilde{\boldsymbol{\theta}}_n$ is somewhere between $\boldsymbol{\theta}_0$ and $\hat{\boldsymbol{\theta}}_n$. If $\frac{1}{n}Dd^2L_n(\tilde{\boldsymbol{\theta}}_n) = O_p(1)$ (bounded in probability), then the third term is $o_p(1)(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0)$ and

$$\sqrt{n}(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0) = \left[-\frac{1}{n}d^2L_n(\boldsymbol{\theta}_0) + o_p(1) \right]^{-1} \frac{\sqrt{n}}{n} \nabla L_n(\boldsymbol{\theta}_0).$$

Now

- * $-\frac{1}{n}d^2L_n(\boldsymbol{\theta}_0) + o_p(1) \rightarrow \mathbf{E}_{\boldsymbol{\theta}_0}[-d^2L(\boldsymbol{\theta}_0)] = \mathbf{I}(\boldsymbol{\theta}_0)$ almost surely by the law of large number.
- * $n^{-1/2}\nabla L_n(\boldsymbol{\theta}_0)$ converges to a multivariate normal with mean $\mathbf{0}_p$ and variance

$$\mathbf{E}_{\boldsymbol{\theta}_0}[\nabla L(\boldsymbol{\theta}_0)dL(\boldsymbol{\theta}_0)],$$

which equals $\mathbf{I}(\boldsymbol{\theta}_0)$ under exchangeability of integral and differentiation.

Then by the Slutsky theorem,

$$\begin{aligned} & \sqrt{n}(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0) \\ \rightarrow & N_p\left(\mathbf{0}_p, \mathbf{I}^{-1}(\boldsymbol{\theta}_0) \cdot \mathbf{E}_{\boldsymbol{\theta}_0}[\nabla \ln f(\boldsymbol{\theta}_0)d \ln f(\boldsymbol{\theta}_0)] \cdot \mathbf{I}^{-1}(\boldsymbol{\theta}_0)\right) \\ = & N_p\left(\mathbf{0}_p, \mathbf{I}^{-1}(\boldsymbol{\theta}_0)\right) \end{aligned}$$

in distribution.

– In practice, we can estimate the variance by

- * Fisher information matrix $\mathbf{I}^{-1}(\hat{\boldsymbol{\theta}})$,
- * observed information matrix $[-(1/n)d^2L_n(\hat{\boldsymbol{\theta}})]^{-1}$, or
- * the sandwich estimator

- Asymptotic efficiency of MLE.

“Cramer-Rao theorem” says the variance of any unbiased estimator is “at least” $(n\mathbf{I}(\boldsymbol{\theta}_0))^{-1}$ (the difference is psd). So MLE has the smallest asymptotic variance within the class of unbiased estimators.

17 Lecture 17, Oct 28

Announcements

- No TA office hours this Friday.
- FAQs on HW6: <http://hua-zhou.github.io/teaching/st758-2014fall/st758fall2014/2014/10/26/hw6-hints.html>
- Roadmap (to winter break!): HW7, HW8, simulation project, final or not?

Last time

- Review of HW5
- MLE

Today

- Newton and scoring algorithm
- Hierarchy of optimization problems

Newton's method and Fisher's scoring (KL Chapter 14)

Joseph Raphson



Consider maximizing log-likelihood $L(\boldsymbol{\theta})$, $\boldsymbol{\theta} \in \Theta \subset \mathbb{R}^p$.

- Newton’s method was originally developed for finding roots of nonlinear equations $f(\boldsymbol{\theta}) = \mathbf{0}$ (KL 5.4).
- Newton’s method (aka Newton-Raphson method) is considered the gold standard for its fast (quadratic) convergence

$$\frac{\|\boldsymbol{\theta}^{(t+1)} - \boldsymbol{\theta}^*\|}{\|\boldsymbol{\theta}^{(t)} - \boldsymbol{\theta}^*\|^2} \rightarrow \text{constant.}$$

- Idea: iterative quadratic approximation.
- Taylor expansion around the current iterate $\boldsymbol{\theta}^{(t)}$

$$L(\boldsymbol{\theta}) \approx L(\boldsymbol{\theta}^{(t)}) + dL(\boldsymbol{\theta}^{(t)})(\boldsymbol{\theta} - \boldsymbol{\theta}^{(t)}) + \frac{1}{2}(\boldsymbol{\theta} - \boldsymbol{\theta}^{(t)})^\top d^2L(\boldsymbol{\theta}^{(t)})(\boldsymbol{\theta} - \boldsymbol{\theta}^{(t)})$$

and then maximize the quadratic approximation.

- To maximize the quadratic function, we equate its gradient to zero

$$\nabla L(\boldsymbol{\theta}^{(t)}) + [d^2L(\boldsymbol{\theta}^{(t)})](\boldsymbol{\theta} - \boldsymbol{\theta}^{(t)}) = \mathbf{0}_p,$$

which suggests the next iterate

$$\begin{aligned} \boldsymbol{\theta}^{(t+1)} &= \boldsymbol{\theta}^{(t)} - [d^2L(\boldsymbol{\theta}^{(t)})]^{-1}\nabla L(\boldsymbol{\theta}^{(t)}) \\ &= \boldsymbol{\theta}^{(t)} + [-d^2L(\boldsymbol{\theta}^{(t)})]^{-1}\nabla L(\boldsymbol{\theta}^{(t)}). \end{aligned}$$

- Some issues with the Newton’s iteration
 - Need to derive, evaluate, and “invert” the observed information matrix. Remedies:
 1. exploit structure in Hessian whenever possible,
 2. numerical differentiation (works for small problems), or
 3. quasi-Newton method (to be discussed later)
 - Stability: Newton’s iterate is not guaranteed to be an ascent algorithm. It’s equally happy to head uphill or downhill. Remedies:
 1. approximate $-d^2L(\boldsymbol{\theta}^{(t)})$ by a positive definite \mathbf{A} (if it’s not), *and*
 2. line search (backtracking).

In summary, Newton's method iterates according to

$$\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} + s[\mathbf{A}^{(t)}]^{-1}\nabla L(\boldsymbol{\theta}^{(t)}) = \boldsymbol{\theta}^{(t)} + s\Delta\boldsymbol{\theta}^{(t)}$$

where $\mathbf{A}^{(t)}$ is a pd approximation of $-d^2L(\boldsymbol{\theta}^{(t)})$ and s is a step length.

Why? By first-order Taylor expansion,

$$\begin{aligned} & L(\boldsymbol{\theta}^{(t)} + s\Delta\boldsymbol{\theta}^{(t)}) - L(\boldsymbol{\theta}^{(t)}) \\ &= dL(\boldsymbol{\theta}^{(t)})s\Delta\boldsymbol{\theta}^{(t)} + o(s) \\ &= sdL(\boldsymbol{\theta}^{(t)})[\mathbf{A}^{(t)}]^{-1}\nabla L(\boldsymbol{\theta}^{(t)}) + o(s). \end{aligned}$$

For s sufficiently small, right hand side is strictly positive.

- Backtracking strategy: step-halving ($s = 1, 1/2, \dots$), golden section search, cubic interpolation, Amijo rule, ...
- How to approximating $-d^2L(\boldsymbol{\theta})$? More of an art than science. Often requires problem specific analysis.
- Taking $\mathbf{A} = \mathbf{I}$ leads to the method of *steepest ascent*, aka *gradient ascent*.
- *Fisher's scoring method*: replace $-d^2L(\boldsymbol{\theta})$ by the expected Fisher information matrix

$$\mathbf{I}(\boldsymbol{\theta}) = \mathbf{E}[-d^2L(\boldsymbol{\theta})] = \mathbf{E}[\nabla L(\boldsymbol{\theta})dL(\boldsymbol{\theta})] \succeq \mathbf{0}_{p \times p},$$

which is psd under exchangeability of expectation and differentiation.

Therefore the Fisher's scoring algorithm iterates according to

$$\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} + s[\mathbf{I}(\boldsymbol{\theta}^{(t)})]^{-1}\nabla L(\boldsymbol{\theta}^{(t)})$$

Hierarchy of optimization problems

Difficulty of optimization problems *in general*

Harder	Easier
discrete (combinatorial) optimization	continuous optimization
non-smooth	smooth
non-convex	convex
constrained	un-constrained
inequality constraint	equality constrained

18 Lecture 18, Oct 30

Announcements

- No TA office hours this Friday. Makeup office hours next Mon, Nov 3 @ 3P-5P.

Last time

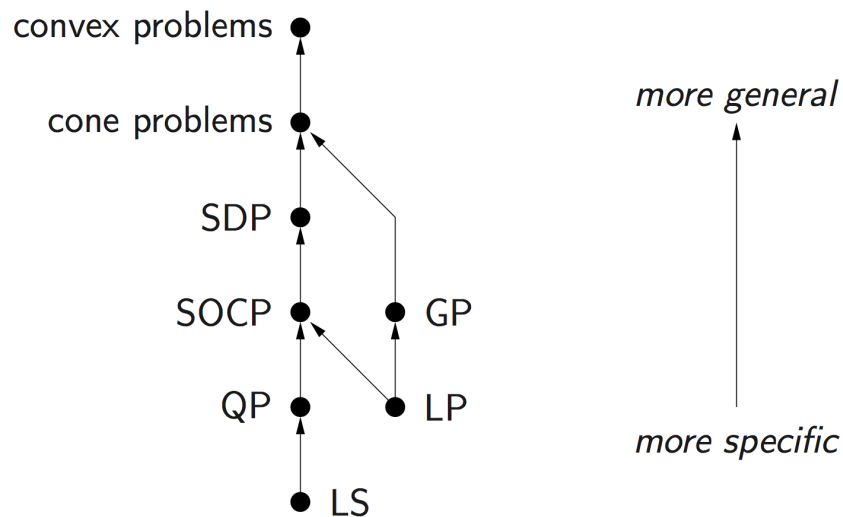
- Newton and Fisher scoring method
- Hierarchy of optimization problems

Today

- Convex optimization
- Some fundamentals of optimization theory


Convex optimization

- *Extremely important* skill to recognize or transform to convex problems



- Examples: ℓ_∞ regression, ℓ_1 regression, quantile regression, and many more.

- *Convex Optimization* by Boyd and Vandenberghe and accompanying slides
<http://www.stanford.edu/~boyd/cvxbook/>



- Lecture videos:
<http://www.stanford.edu/class/ee364a/videos.html>
<http://www.stanford.edu/class/ee364b/videos.html>
- Convex programming (LS, LP, QP, GP, SOCP, SDP) is almost becoming a technology (Cplex, Gurobi, Mosek, cvx, Matlab, ...)
- Non-convex optimization still occurs in many natural statistical applications. Statisticians have specialized tools to deal with them (Fisher scoring method, EM algorithm, simulated annealing, ...)

Unconstrained optimization (KL 11.2)




Figure 1 Unconstrained optimization in one variable

- Possible confusion:
 - We (statisticians) talk about *maximization*: $\max L_n(\boldsymbol{\theta})$.
 - People talk about *minimization* in the optimization world: $\min_{\boldsymbol{x}} f(\boldsymbol{x})$.
- Fundamental questions: When does a function have minimum? How do we tell whether a point is minimum?
- When does a function have a minimum?
(Weierstrass) A continuous function $f(x)$ defined on a compact (closed and bounded) set is bounded below and attains its minimum.
- None of the Weierstrass conditions can be taken out.
 - $f(x) = x, x \in (-\infty, \infty)$. Non-compact support.
 - $f(x) = \tan(x), x \in (-\pi/2, \pi/2)$. Non-compact support.
 - $f(x) = x, x \in (-1, 1)$ and $f(-1) = f(1) = 0$. The minimum not attained by the *discontinuous* function f .
- None of the Weierstrass conditions are necessary. $f(x) = x, x \in [0, 2), f(x) = 1, x \in (2, \infty)$.

- Coercive function: $\{\mathbf{x} \in U : f(\mathbf{x}) \leq f(\mathbf{y})\}$ is compact for all $\mathbf{y} \in U$.
Weierstrass theorem also holds for a coercive function defined on a possibly open U .
- Necessary conditions for a local minimum.
Assume f has a local minimum at interior point $\mathbf{y} \in U$.
 - (Fermat) If f is differentiable, then $\nabla f(\mathbf{x})$ vanishes at \mathbf{y} .
 - If f is twice differentiable, then $d^2 f(\mathbf{y})$ is psd.
- Points with $\nabla f(\mathbf{x}) = \mathbf{0}$ are called *stationary points* or *critical points*. Most optimization algorithms try to find the stationary points of the function and then check sufficient condition.
- Counter-examples to necessary conditions. (1) $f(x) = x^3$ has zero gradient at 0, which is not local minimum. (2) $f(x) = |x|$ has local minimum at 0, where the gradient does not exist.
- (A first-order sufficient condition; first derivative test) Suppose f is differentiable in a ball $B(\mathbf{y})$ around an interior point \mathbf{y} , and $\langle \nabla f(\mathbf{x}), \mathbf{x} - \mathbf{y} \rangle \geq 0$ for all $\mathbf{x} \in B(\mathbf{y})$, then \mathbf{y} is a local minimum.
- (A second-order sufficient condition; second derivative test) If $\nabla f(\mathbf{y}) = \mathbf{0}$ and $d^2 f(\mathbf{y})$ is pd, then \mathbf{y} is a strict local minimum.
- Remark: In case $d^2 f(\mathbf{y})$ is neither positive definite nor negative definite but non-singular, \mathbf{y} is a *saddle point*, i.e., a stationary point that is neither a local minimum nor a local maximum. In case $d^2 f(\mathbf{y})$ is singular, we cannot tell.
- Example: $f_1(x, y) = x^4 + y^4$, $f_2(x, y) = -x^4 - y^4$, $f_3(x, y) = x^3 + y^3$. Origin is a stationary (critical) point and the Hessian $d^2 f_i(0, 0) = \mathbf{0}_{2 \times 2}$ is singular. Origin is a minimum, maximum, and a saddle point respectively.

Convexity and global optima




FIGURE 11.2. Plot of the Convex Function $e^x + x^2$

- $f : U \mapsto \mathbb{R}$ is *convex* if
 - U is a convex set ($\lambda \mathbf{x} + (1 - \lambda)\mathbf{y} \in U$ for all $\mathbf{x}, \mathbf{y} \in U$ and $\lambda \in (0, 1)$), and
 - $f(\lambda \mathbf{x} + (1 - \lambda)\mathbf{y}) \leq \lambda f(\mathbf{x}) + (1 - \lambda)f(\mathbf{y})$, for all $\mathbf{x}, \mathbf{y} \in U$ and $\lambda \in (0, 1)$.

f is *strictly convex* if the inequality is strict for all $\mathbf{x} \neq \mathbf{y} \in U$ and λ .
- (*Supporting hyperplane inequality*) A differentiable function f is convex if and only if $f(\mathbf{x}) \geq f(\mathbf{y}) + \langle \nabla f(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle$ for all $\mathbf{x}, \mathbf{y} \in U$.
- (*Second-order condition for convexity*) A twice differentiable function f is convex if and only if $d^2f(\mathbf{x})$ is psd for all $\mathbf{x} \in U$.
It is strictly convex if $d^2f(\mathbf{x})$ is pd for all $\mathbf{x} \in U$.
- (*Convexity and global optima*) Suppose f is a convex function on a convex set U .
 1. Any stationary point \mathbf{y} is a global minimum. (By supporting hyperplane inequality, $f(\mathbf{x}) \geq f(\mathbf{y}) + \langle \nabla f(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle = f(\mathbf{y})$ for all $\mathbf{x} \in U$.)
 2. Any local minimum is a global minimum.

3. The set of (global) minima $\{\mathbf{x} \in U : f(\mathbf{x}) = f(\mathbf{y})\}$ is convex.
4. If f is strictly convex, then the global minimum, if exists, is unique.
- Example: Least squares estimate. $f(\boldsymbol{\beta}) = \frac{1}{2}\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2$ has Hessian $d^2f = \mathbf{X}^\top \mathbf{X}$ which is psd. So f is convex and any stationary point (solution to the normal equation) is a global minimum. When \mathbf{X} is rank deficient, the set of solutions is convex.
 - (Jensen's inequality) W a random variable taking values in U and h is convex on U . Then

$$\mathbf{E}[h(W)] \geq h[\mathbf{E}(W)],$$

provided both expectations exist. For a strictly convex h , equality holds if and only if $W = \mathbf{E}(W)$ almost surely.

Proof: supporting hyperplane inequality taking $\mathbf{x} = \mathbf{W}$ and $\mathbf{y} = \mathbf{E}(\mathbf{W})$.

- (Information inequality) Let f and g be two densities with respect to a common measure μ . $h, g > 0$ almost everywhere relative to μ . Then

$$\mathbf{E}_f(\ln f) \geq \mathbf{E}_f(\ln g),$$

with equality if and only if $f = g$ almost everywhere on μ .

Proof: Apply Jensen's inequality to the convex function $-\ln(t)$ and random variable $W = g(x)/f(x)$.

Applications: M-estimation, EM algorithm.

19 Lecture 19, Nov 4

Announcements

- Answer to Caleb's question: Why do people care about ℓ_∞ regression?

$$\min_{\boldsymbol{\beta}} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_\infty = \min_{\boldsymbol{\beta}} \max_i |y_i - \mathbf{x}_i^T \boldsymbol{\beta}|.$$

Chebyshev approximation, or minimax approximation

$$f(x) \sim \sum_{i=0}^N c_i T_i(x).$$

For fixed degree N , find the coefficients c_i that minimize the worst possible approximation error.

Last time

- Convex optimization
- Fundamentals of optimization theory: optimality conditions for unconstrained optimization, convexity and global optima

Today

- Fundamentals of optimization theory (cont'd): optimality conditions for constrained optimization
- Application of Newton and Fisher scoring algorithm: GLM

Optimization with equality constraints (KL 11.3)

Consider the equality constrained minimization problem

$$\begin{aligned} & \text{minimize} && f(\mathbf{x}) \\ & \text{subject to} && g_i(\mathbf{x}) = 0, i = 1, \dots, m \\ & && \mathbf{x} \in U \subset \mathbb{R}^n. \end{aligned}$$

We write

$$g(\mathbf{x}) = \begin{pmatrix} g_1(\mathbf{x}) \\ \vdots \\ g_m(\mathbf{x}) \end{pmatrix} \in \mathbb{R}^m \text{ and } Dg(\mathbf{x}) = \begin{pmatrix} dg_1(\mathbf{x}) \\ \dots \\ dg_m(\mathbf{x}) \end{pmatrix} \in \mathbb{R}^{m \times n}.$$

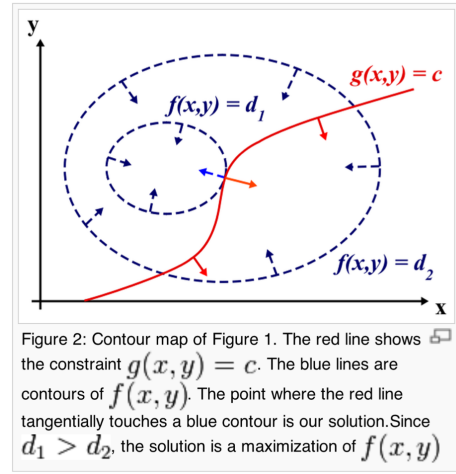
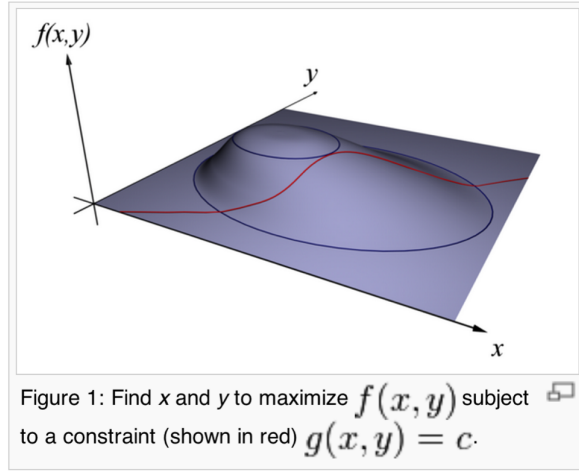
- Method of Lagrange multiplier. *Lagrangian* function

$$L(\mathbf{x}, \boldsymbol{\lambda}) = f(\mathbf{x}) + \boldsymbol{\lambda}^\top g(\mathbf{x}) = f(\mathbf{x}) + \sum_{i=1}^m \lambda_i g_i(\mathbf{x}).$$

Strategy for finding the equality constrained minimum: find the stationary point $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ of the Lagrangian,

$$\begin{aligned} \nabla_{\mathbf{x}} L(\mathbf{x}, \boldsymbol{\lambda}) &= \nabla f(\mathbf{x}) + \sum_{i=1}^m \lambda_i \nabla g_i(\mathbf{x}) = \mathbf{0}_n \\ g(\mathbf{x}) &= \mathbf{0}_m. \end{aligned}$$

- Intuition: Null space of the matrix Dg is the tangent space. Movement along the tangent space does not change constraint function values. We need the ∇f to be orthogonal to the tangent space. In other words, ∇f is in the column space of $[Dg]^T$.
- Intuition of the Lagrange multiplier method: Hill climb along a trail which is a contour line of the constraint function. We feel effortless exactly when the direction of our movement is perpendicular to the steepest ascent direction of the hill. In other words, steepest ascent direction of the constraint function aligns with that of the hill.



- (Necessary condition for a constrained local minimum) Assume conditions (i) $g(\mathbf{y}) = \mathbf{0}_m$, (2) f and g are differentiable in some n -ball $B(\mathbf{y})$, (iii) $Dg(\mathbf{y}) \in \mathbb{R}^{m \times n}$ is continuous at \mathbf{y} , (iv) $Dg(\mathbf{y})$ has full row rank, (v) $f(\mathbf{x}) \geq f(\mathbf{y})$ for any $\mathbf{x} \in B(\mathbf{y})$ satisfying $g(\mathbf{x}) = \mathbf{0}_m$ (\mathbf{y} a local minimum subject to constraints). Then there exists $\boldsymbol{\lambda} \in \mathbb{R}^m$ satisfying $\nabla f(\mathbf{y}) + \sum_{i=1}^m \lambda_i \nabla g_i(\mathbf{y}) = \mathbf{0}_n$, i.e., $(\mathbf{y}, \boldsymbol{\lambda})$ is a stationarity point of the Lagrangian $L(\mathbf{x}, \boldsymbol{\lambda})$. In other words, there exists $\boldsymbol{\lambda} \in \mathbb{R}^m$, such that $\nabla L(\mathbf{y}, \boldsymbol{\lambda}) = \mathbf{0}_{m+n}$.
- (Sufficient condition for a constrained local minimum) (i) f twice differentiable at \mathbf{y} , (ii) g twice differentiable at \mathbf{y} , (iii) the Jacobian matrix $Dg(\mathbf{y}) \in \mathbb{R}^{m \times n}$ has full row rank m , (iv) it is a stationarity point of the Lagrangian at a given $\boldsymbol{\lambda} \in \mathbb{R}^m$, (v) $\mathbf{u}^\top d^2 f(\mathbf{y}) \mathbf{u} > 0$ for all $\mathbf{u} \neq \mathbf{0}_n$ satisfying $[Dg(\mathbf{y})] \mathbf{u} = \mathbf{0}_m$ (tangent vectors). Then \mathbf{y} is a strict local minimum of f under constraint $g(\mathbf{y}) = \mathbf{0}_m$.
- Check condition (v). Condition (v) is equivalent to the “bordered determinantal criterion”

$$(-1)^m \det \begin{pmatrix} \mathbf{0}_{m \times m} & \mathbf{B}_r \\ \mathbf{B}_r^\top & \mathbf{A}_{rr} \end{pmatrix} > 0$$

for $r = m + 1, \dots, n$, where

- \mathbf{A}_{rr} is the top left r -by- r block of $d^2 f(\mathbf{y}) + \sum_{i=1}^m \lambda_i d^2 g_i(\mathbf{y})$
- $\mathbf{B}_r \in \mathbb{R}^{m \times r}$ is the first r columns of the $Dg(\mathbf{y})$.

- (Sufficient condition for a global constrained minimum) Lagrangian first order condition + convexity of the Lagrangian on U .
- (Interpretation of the Lagrange multipliers $\boldsymbol{\lambda}$). Consider $\min f(\mathbf{x})$ subject to some resource constraint $g(\mathbf{x}) = \mathbf{b}$. Consider the solution $\mathbf{x}^*(\mathbf{b})$ as a function of \mathbf{b} . Then it can be shown that

$$\frac{\partial f(\mathbf{x}^*(\mathbf{b}))}{\partial b_j} = \lambda_j.$$

That's why the score test is classically called the Lagrange multiplier test.

- Example: Linearly constrained least squares solution. $\min \frac{1}{2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2$ subject to linear constrained $\mathbf{V}\boldsymbol{\beta} = \mathbf{d}$. Form the Lagrangian

$$L(\boldsymbol{\beta}, \boldsymbol{\lambda}) = \frac{1}{2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \boldsymbol{\lambda}^\top (\mathbf{V}\boldsymbol{\beta} - \mathbf{d}).$$

Stationary condition says

$$\begin{aligned} \mathbf{X}^\top \mathbf{X}\boldsymbol{\beta} - \mathbf{X}^\top \mathbf{y} + \mathbf{V}^\top \boldsymbol{\lambda} &= \mathbf{0}_p \\ \mathbf{V}\boldsymbol{\beta} &= \mathbf{d} \end{aligned}$$

or equivalently

$$\begin{pmatrix} \mathbf{X}^\top \mathbf{X} & \mathbf{V}^\top \\ \mathbf{V} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \boldsymbol{\beta} \\ \boldsymbol{\lambda} \end{pmatrix} = \begin{pmatrix} \mathbf{X}^\top \mathbf{y} \\ \mathbf{d} \end{pmatrix},$$

which can be solved by say sweeping on

$$\begin{pmatrix} \mathbf{X}^\top \mathbf{X} & \mathbf{V}^\top & \mathbf{X}^\top \mathbf{y} \\ \mathbf{V} & \mathbf{0} & \mathbf{d} \\ \mathbf{y}\mathbf{X}^\top & \mathbf{d}^\top & \mathbf{y}^\top \mathbf{y} \end{pmatrix},$$

or Cholesky or QR.

Optimization with both equality and inequality constraints (KL 11.4)

Consider the constrained minimization problem

$$\begin{aligned} &\text{minimize} && f(\mathbf{x}) \\ &\text{subject to} && g_i(\mathbf{x}) = 0, i = 1, \dots, p \\ &&& h_j(\mathbf{x}) \leq 0, i = 1, \dots, q \\ &&& \mathbf{x} \in U \subset \mathbb{R}^n. \end{aligned}$$

- Lagrangian function:

$$L(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = f(\mathbf{x}) + \sum_{i=1}^p \lambda_i g_i(\mathbf{x}) + \sum_{j=1}^q \mu_j h_j(\mathbf{x}).$$



- Karush-Kuhn-Tucker (KKT) necessary condition: If (1) \mathbf{y} is a local constrained minimum and (2) satisfies certain constraint qualifications (Kuhn-Tucker, Mangasarian-Fromovitz), then

1. (Lagrangian stationarity condition) there exist $\boldsymbol{\lambda} \in \mathbb{R}^p$, $\boldsymbol{\mu} \in \mathbb{R}^q$ such that

$$\nabla f(\mathbf{y}) + \sum_{i=1}^p \lambda_i \nabla g_i(\mathbf{y}) + \sum_{j=1}^q \mu_j \nabla h_j(\mathbf{y}) = \mathbf{0},$$

2. (Complementary slackness) $\mu_j = 0$ if $h_j(\mathbf{y}) < 0$ and $\mu_j > 0$ otherwise.

- Sufficient condition: KKT + second order condition.
- Global minimum: KKT conditions + convexity.
- Read KL Section 11.4 for more details. KKT is “one of the great triumphs of 20th century applied mathematics”.

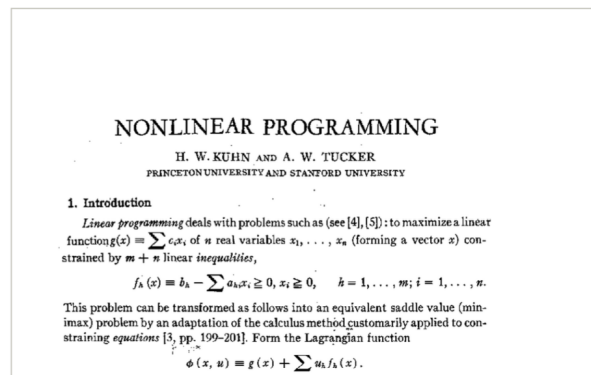
1. [^] Kuhn, H. W.; Tucker, A. W. (1951). "Nonlinear programming" . *Proceedings of 2nd Berkeley Symposium*. Berkeley: University of California Press. pp. 481–492. [MR47303](#) 
2. [^] W. Karush (1939). *Minima of Functions of Several Variables with Inequalities as Side Constraints*. M.Sc. Dissertation. Dept. of Mathematics, Univ. of Chicago, Chicago, Illinois.

Nonlinear Programming

H. W. Kuhn, and A. W. Tucker

Source: Proc. Second Berkeley Symp. on Math. Statist. and Prob. (Univ. of Calif. Press, 1951), 481-492.

First Page: [Hide](#)



Generalized linear model (GLM) (KL 14.7, JM 9.7)

Let's consider a concrete example: logistic regression.

- You want to make a fame by participating the Capital One competition. The goal is to predict whether a credit card transaction is fraud ($y_i = 1$) or not ($y_i = 0$). Predictors (\mathbf{x}_i) include: time of transaction, last location, merchant, ...
- $y_i \in \{0, 1\}$, $\mathbf{x}_i \in \mathbb{R}^p$. Model $y_i \sim \text{Bernoulli}(p_i)$.

- Logistic regression. Density

$$\begin{aligned}
f(y_i|p_i) &= p_i^{y_i}(1-p_i)^{1-y_i} \\
&= e^{y_i \ln p_i + (1-y_i) \ln(1-p_i)} \\
&= e^{y_i \ln \frac{p_i}{1-p_i} + \ln(1-p_i)},
\end{aligned}$$

where

$$\begin{aligned}
E(y_i) = p_i &= \frac{e^{\mathbf{x}_i^\top \boldsymbol{\beta}}}{1 + e^{\mathbf{x}_i^\top \boldsymbol{\beta}}} \quad (\text{mean function, inverse link function}) \\
\mathbf{x}_i^\top \boldsymbol{\beta} &= \ln \left(\frac{p_i}{1-p_i} \right) \quad (\text{logit link function}).
\end{aligned}$$

- Given data (y_i, \mathbf{x}_i) , $i = 1, \dots, n$,

$$\begin{aligned}
L_n(\boldsymbol{\beta}) &= \sum_{i=1}^n [y_i \ln p_i + (1-y_i) \ln(1-p_i)] \\
&= \sum_{i=1}^n \left[y_i \mathbf{x}_i^\top \boldsymbol{\beta} - \ln(1 + e^{\mathbf{x}_i^\top \boldsymbol{\beta}}) \right] \\
\nabla L_n(\boldsymbol{\beta}) &= \sum_{i=1}^n \left(y_i \mathbf{x}_i - \frac{e^{\mathbf{x}_i^\top \boldsymbol{\beta}}}{1 + e^{\mathbf{x}_i^\top \boldsymbol{\beta}}} \mathbf{x}_i \right) \\
&= \sum_{i=1}^n (y_i - p_i) \mathbf{x}_i = \mathbf{X}^\top (\mathbf{y} - \mathbf{p}) \\
-d^2 L_n(\boldsymbol{\beta}) &= \sum_{i=1}^n p_i (1-p_i) \mathbf{x}_i \mathbf{x}_i^\top = \mathbf{X}^\top \mathbf{W} \mathbf{X}, \\
&\quad \text{where } \mathbf{W} = \text{diag}(w_1, \dots, w_n), w_i = p_i(1-p_i) \\
\mathbf{I}_n(\boldsymbol{\beta}) &= \mathbf{E}[-d^2 L_n(\boldsymbol{\beta})] = -d^2 L_n(\boldsymbol{\beta}).
\end{aligned}$$

- Newton's method = Fisher's scoring iteration:

$$\begin{aligned}
\boldsymbol{\beta}^{(t+1)} &= \boldsymbol{\beta}^{(t)} + s[-d^2 L(\boldsymbol{\beta}^{(t)})]^{-1} \nabla L(\boldsymbol{\beta}^{(t)}) \\
&= \boldsymbol{\beta}^{(t)} + s(\mathbf{X}^\top \mathbf{W}^{(t)} \mathbf{X})^{-1} \mathbf{X}^\top (\mathbf{y} - \mathbf{p}^{(t)}) \\
&= (\mathbf{X}^\top \mathbf{W}^{(t)} \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{W}^{(t)} \left[\mathbf{X} \boldsymbol{\beta}^{(t)} + s(\mathbf{W}^{(t)})^{-1} (\mathbf{y} - \mathbf{p}^{(t)}) \right] \\
&= (\mathbf{X}^\top \mathbf{W}^{(t)} \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{W}^{(t)} \mathbf{z}^{(t)},
\end{aligned}$$

where

$$\mathbf{z}^{(t)} = \mathbf{X} \boldsymbol{\beta}^{(t)} + s(\mathbf{W}^{(t)})^{-1} (\mathbf{y} - \mathbf{p}^{(t)})$$

are the working responses. A Newton's iteration is equivalent to solving a weighed least squares problem $\sum_{i=1}^n w_i(z_i - \mathbf{x}_i^T \boldsymbol{\beta})^2$. Thus the name IRWLS (iteratively re-weighted least squares).

Common distributions with typical uses and canonical link functions

Distribution	Support of distribution	Typical uses	Link name	Link function	Mean function
Normal	real: $(-\infty, +\infty)$	Linear-response data	Identity	$\mathbf{X}\boldsymbol{\beta} = \mu$	$\mu = \mathbf{X}\boldsymbol{\beta}$
Exponential	real: $(0, +\infty)$	Exponential-response data, scale parameters	Inverse	$\mathbf{X}\boldsymbol{\beta} = \mu^{-1}$	$\mu = (\mathbf{X}\boldsymbol{\beta})^{-1}$
Gamma					
Inverse Gaussian	real: $(0, +\infty)$		Inverse squared	$\mathbf{X}\boldsymbol{\beta} = \mu^{-2}$	$\mu = (\mathbf{X}\boldsymbol{\beta})^{-1/2}$
Poisson	integer: $[0, +\infty)$	count of occurrences in fixed amount of time/space	Log	$\mathbf{X}\boldsymbol{\beta} = \ln(\mu)$	$\mu = \exp(\mathbf{X}\boldsymbol{\beta})$
Bernoulli	integer: $[0, 1]$	outcome of single yes/no occurrence	Logit	$\mathbf{X}\boldsymbol{\beta} = \ln\left(\frac{\mu}{1-\mu}\right)$	$\mu = \frac{\exp(\mathbf{X}\boldsymbol{\beta})}{1 + \exp(\mathbf{X}\boldsymbol{\beta})} = \frac{1}{1 + \exp(-\mathbf{X}\boldsymbol{\beta})}$
Binomial	integer: $[0, N]$	count of # of "yes" occurrences out of N yes/no occurrences			
Categorical	integer: $[0, K)$	outcome of single K-way occurrence			
	K-vector of integer: $[0, 1]$, where exactly one element in the vector has the value 1				
Multinomial	K-vector of integer: $[0, N]$	count of occurrences of different types (1 .. K) out of N total K-way occurrences			

Let's consider the more general class of generalized linear models (GLM).

- Y belongs to an exponential family with density

$$p(y|\theta, \phi) = \exp\left\{\frac{y\theta - b(\theta)}{a(\phi)} + c(y, \phi)\right\}.$$

θ : natural parameter. $\phi > 0$: dispersion parameter. GLM relates the mean $\mu = \mathbf{E}(Y|\mathbf{x})$ via a strictly increasing link function

$$g(\mu) = \eta = \mathbf{x}^T \boldsymbol{\beta}, \quad \mu = g^{-1}(\eta)$$

- Score, Hessian, information

$$\begin{aligned}\nabla L_n(\boldsymbol{\beta}) &= \sum_{i=1}^n \frac{(y_i - \mu_i)\mu'_i(\eta_i)}{\sigma_i^2} \mathbf{x}_i \\ -d^2 L_n(\boldsymbol{\beta}) &= \sum_{i=1}^n \frac{[\mu'_i(\eta_i)]^2}{\sigma_i^2} \mathbf{x}_i \mathbf{x}_i^\top - \sum_{i=1}^n \frac{(y_i - \mu_i)\theta''(\eta_i)}{\sigma_i^2} \mathbf{x}_i \mathbf{x}_i^\top \\ \mathbf{I}_n(\boldsymbol{\beta}) &= \mathbf{E}[-d^2 L_n(\boldsymbol{\beta})] = \sum_{i=1}^n \frac{[\mu'_i(\eta_i)]^2}{\sigma_i^2} \mathbf{x}_i \mathbf{x}_i^\top = \mathbf{X}^\top \mathbf{W} \mathbf{X}.\end{aligned}$$

- Fisher scoring method

$$\boldsymbol{\beta}^{(t+1)} = \boldsymbol{\beta}^{(t)} + s[\mathbf{I}(\boldsymbol{\beta}^{(t)})]^{-1} \nabla L_n(\boldsymbol{\beta}^{(t)})$$

IRWLS with weights $w_i = [\mu'_i(\eta_i)]^2 / \sigma_i^2$ and some working responses z_i .

- For *canonical link*, $\theta = \eta$, the second term of Hessian vanishes and Hessian coincides with Fisher information matrix. Convex problem ☺

Fisher's scoring = Newton's method.

- Non-canonical link, non-convex problem ☹

Fisher's scoring algorithm \neq Newton's method.

Example: Probit regression (binary response with probit link). $y_i \sim \text{Bernoulli}(p_i)$ and

$$p_i = \Phi(\mathbf{x}_i^T \boldsymbol{\beta}), \quad \eta_i = \mathbf{x}_i^T \boldsymbol{\beta} = \Phi^{-1}(p_i),$$

where $\Phi(\cdot)$ is the cdf of a standard normal.

- `glmfit()` in R and MATLAB implements the Fisher scoring method, aka IRWLS, for GLMs.

20 Lecture 20, Nov 6

Announcements

- HW6 due today. Submit hardcopy and email code to me (LastFirstHW6.R or LastFirstHW6.Rmd)
- HW7 due date changed to Tue Nov 18

Last time

- Optimality conditions for constrained optimization (KKT)
- Application of Newton and Fisher scoring algorithm: GLM

Today

- Application of Newton and Fisher scoring algorithm: nonlinear regression
- EM algorithm


Nonlinear regression – Gauss-Newton method (KL 14.4-14.6, JM 9.8)

- Now we finally get to the problem Gauss faced in 1800!
Relocate Ceres by fitting 41 observations to a 6-parameter (nonlinear) orbit.
- Nonlinear least squares (curve fitting):

$$\text{minimize } f(\boldsymbol{\beta}) = \frac{1}{2} \sum_{i=1}^n [y_i - \mu_i(\mathbf{x}_i, \boldsymbol{\beta})]^2$$

For example, y_i = dry weight of onion and x_i = growth time, and we want to fit a 3-parameter growth curve

$$\mu(x, \beta_1, \beta_2, \beta_3) = \frac{\beta_3}{1 + e^{-\beta_1 - \beta_2 x}}.$$



- “Score” and “information matrices”

$$\begin{aligned} \nabla f(\boldsymbol{\beta}) &= -\sum_{i=1}^n [y_i - \mu_i(\boldsymbol{\beta})] \nabla \mu_i(\boldsymbol{\beta}) \\ d^2 f(\boldsymbol{\beta}) &= \sum_{i=1}^n \nabla \mu_i(\boldsymbol{\beta}) d\mu_i(\boldsymbol{\beta}) - \sum_{i=1}^n [y_i - \mu_i(\boldsymbol{\beta})] d^2 \mu_i(\boldsymbol{\beta}) \\ \mathbf{I}(\boldsymbol{\beta}) &= \sum_{i=1}^n \nabla \mu_i(\boldsymbol{\beta}) d\mu_i(\boldsymbol{\beta}) = \mathbf{J}(\boldsymbol{\beta})^\top \mathbf{J}(\boldsymbol{\beta}), \end{aligned}$$

where $\mathbf{J}(\boldsymbol{\beta})^\top = [\nabla \mu_1(\boldsymbol{\beta}), \dots, \nabla \mu_n(\boldsymbol{\beta})] \in \mathbb{R}^{p \times n}$.

- *Gauss-Newton* (= “Fisher’s scoring algorithm”) uses $\mathbf{I}(\boldsymbol{\beta})$, which is always psd.

$$\boxed{\boldsymbol{\beta}^{(t+1)} = \boldsymbol{\beta}^{(t)} + s \mathbf{I}(\boldsymbol{\beta}^{(t)})^{-1} \nabla L(\boldsymbol{\beta}^{(t)})}$$

- *Levenberg-Marquardt* method, aka *damped least squares algorithm* (DLS), adds a ridge term to the approximate Hessian

$$\boxed{\boldsymbol{\beta}^{(t+1)} = \boldsymbol{\beta}^{(t)} + s [\mathbf{I}(\boldsymbol{\beta}^{(t)}) + \tau \mathbf{I}_p]^{-1} \nabla L(\boldsymbol{\beta}^{(t)})}$$

bridging between Gauss-Newton and steepest descent.

- Other approximation to Hessians: nonlinear GLMs.
See KL 14.4 for examples.

Which statistical papers are most cited?

Paper	Citations	Per Year
Kaplan-Meier (Kaplan and Meier, 1958)	44502	795
EM (Dempster et al., 1977a)	39167	1059
Cox model (Cox, 1972)	38036	906
Metropolis (Metropolis et al., 1953)	28093	461
FDR (Benjamini and Hochberg, 1995)	24093	1268
Unit root test (Dickey and Fuller, 1979)	15354	439
Lasso (Tibshirani, 1996)	11438	635
bootstrap (Efron, 1979)	10953	313
FFT (Cooley and Tukey, 1965)	10134	207
Gibbs sampler (Gelfand and Smith, 1990)	5901	246

- Citation counts from Google Scholar on 11/03/2014.
- EM is one of the most influential statistical ideas, finding applications in various branches of science.

EM algorithm

- History: Dempster et al. (1977b).

[\[PDF\] Maximum likelihood from incomplete data via the EM algorithm](#)
 AP Dempster, NM Laird, DB Rubin - *Journal of the Royal Statistical Society. ...*, 1977 - JSTOR
 A broadly applicable **algorithm** for computing **maximum likelihood** estimates from **incomplete data** is presented at various levels of generality. Theory showing the monotone behaviour of the **likelihood** and convergence of the **algorithm** is derived. Many examples are sketched, ...
 Cited by 39167 Related articles All 76 versions Web of Science: 16067 Cite Save More

Same idea appears in parameter estimation in HMM (Baum-Welch algorithm) (Baum et al., 1970).

[A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains](#)
 LE Baum, T Petrie, G Soules, N Weiss - *The annals of mathematical statistics*, 1970 - JSTOR
 PYI... YT (**A**, **a**, **f**) and the difficult analysis of **maximizing** this function of **A** for very special choices of **f** presented in [2],[8] that a simple explicit procedure for **maximization** for a general **f** would be quite difficult; however, this is not the case.
 Cited by 3102 Related articles All 4 versions Cite

- Notations

- \mathbf{Y} : observed data
- \mathbf{Z} : missing data
- $\mathbf{X} = (\mathbf{Y}, \mathbf{Z})$: complete data
- Goal: maximize the log-likelihood of the observed data $\ln g(\mathbf{y}|\boldsymbol{\theta})$ (optimization!)
- Idea: choose \mathbf{Z} such that MLE for the complete data is trivial.
- Let $f(\mathbf{x}|\boldsymbol{\theta}) = f(\mathbf{y}, \mathbf{z}|\boldsymbol{\theta})$ be the density of complete data
- Iterative two step procedure

- E step: calculate the conditional expectation

$$Q(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)}) = \mathbf{E}_{\mathbf{Z}|\mathbf{Y}=\mathbf{y},\boldsymbol{\theta}^{(t)}}[\ln f(\mathbf{Y}, \mathbf{Z}|\boldsymbol{\theta}) | \mathbf{Y} = \mathbf{y}, \boldsymbol{\theta}^{(t)}]$$

- M step: maximize $Q(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)})$ to generate the next iterate

$$\boldsymbol{\theta}^{(t+1)} = \operatorname{argmax}_{\boldsymbol{\theta}} Q(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)})$$

- (Ascent property of EM algorithm) By the information inequality,

$$\begin{aligned} & Q(\boldsymbol{\theta} | \boldsymbol{\theta}^{(t)}) - \ln g(\mathbf{y}|\boldsymbol{\theta}) \\ &= \mathbf{E}[\ln f(\mathbf{Y}, \mathbf{Z}|\boldsymbol{\theta})|\mathbf{Y} = \mathbf{y}, \boldsymbol{\theta}^{(t)}] - \ln g(\mathbf{y}|\boldsymbol{\theta}) \\ &= \mathbf{E} \left\{ \ln \left[\frac{f(\mathbf{Y}, \mathbf{Z} | \boldsymbol{\theta})}{g(\mathbf{Y} | \boldsymbol{\theta})} \right] \mid \mathbf{Y} = \mathbf{y}, \boldsymbol{\theta}^{(t)} \right\} \\ &\leq \mathbf{E} \left\{ \ln \left[\frac{f(\mathbf{Y}, \mathbf{Z} | \boldsymbol{\theta}^{(t)})}{g(\mathbf{Y} | \boldsymbol{\theta}^{(t)})} \right] \mid \mathbf{Y} = \mathbf{y}, \boldsymbol{\theta}^{(t)} \right\} \\ &= Q(\boldsymbol{\theta}^{(t)} | \boldsymbol{\theta}^{(t)}) - \ln g(\mathbf{y}|\boldsymbol{\theta}^{(t)}). \end{aligned}$$

Rearranging shows that (fundamental inequality of EM)

$$\ln g(\mathbf{y} | \boldsymbol{\theta}) \geq Q(\boldsymbol{\theta} | \boldsymbol{\theta}^{(t)}) - Q(\boldsymbol{\theta}^{(t)} | \boldsymbol{\theta}^{(t)}) + \ln g(\mathbf{y} | \boldsymbol{\theta}^{(t)})$$

for all $\boldsymbol{\theta}$ and in particular

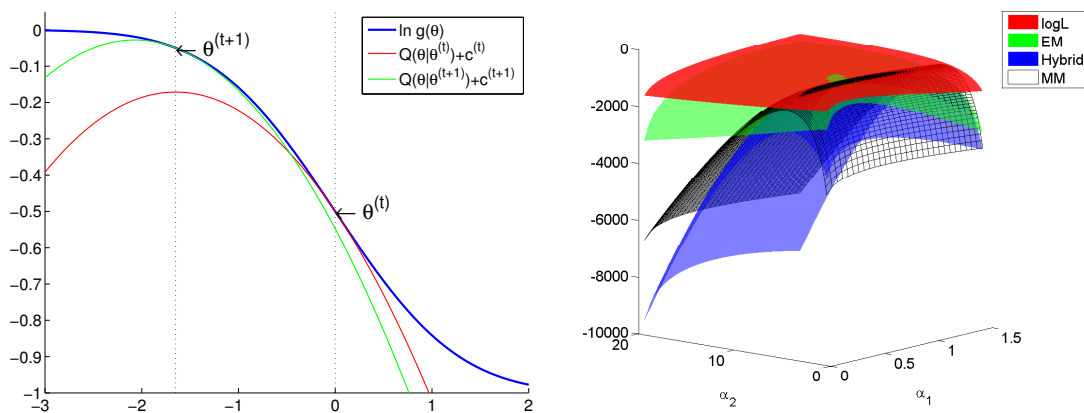
$$\begin{aligned} \ln g(\mathbf{y} | \boldsymbol{\theta}^{(t+1)}) &\geq Q(\boldsymbol{\theta}^{(t+1)} | \boldsymbol{\theta}^{(t)}) - Q(\boldsymbol{\theta}^{(t)} | \boldsymbol{\theta}^{(t)}) + \ln g(\mathbf{y} | \boldsymbol{\theta}^{(t)}) \\ &\geq \ln g(\mathbf{y} | \boldsymbol{\theta}^{(t)}). \end{aligned}$$

Obviously we only need

$$Q(\boldsymbol{\theta}^{(t+1)} | \boldsymbol{\theta}^{(t)}) - Q(\boldsymbol{\theta}^{(t)} | \boldsymbol{\theta}^{(t)}) \geq 0$$


for this ascent property to hold (*generalized EM*).

- Intuition? Keep these pictures in mind



- Under mild regularity conditions, $\boldsymbol{\theta}^{(t)}$ converges to a stationary point of $\ln g(\mathbf{y}|\boldsymbol{\theta})$.
- Numerous applications of EM:
 - finite mixture model, HMM (Baum-Welch algorithm), factor analysis, variance components model aka linear mixed model (LMM), hyper-parameter estimation in empirical Bayes procedure $\max_{\boldsymbol{\alpha}} \int f(\mathbf{y}|\boldsymbol{\theta})\pi(\boldsymbol{\theta}|\boldsymbol{\alpha}) d\boldsymbol{\theta}$ (e.g., HW6/7), missing data, group/censored/truncated model, ...

A canonical example: finite mixture models



- Gaussian finite mixture models: mixture density

$$h(\mathbf{y}) = \sum_{j=1}^k \pi_j h_j(\mathbf{y} \mid \boldsymbol{\mu}_j, \boldsymbol{\Omega}_j), \quad \mathbf{y} \in \mathbb{R}^d,$$

where

$$h_j(\mathbf{y} \mid \boldsymbol{\mu}_j, \boldsymbol{\Omega}_j) = \left(\frac{1}{2\pi}\right)^{d/2} |\det(\boldsymbol{\Omega}_j)|^{-1/2} e^{-\frac{1}{2}(\mathbf{y}-\boldsymbol{\mu}_j)^\top \boldsymbol{\Omega}_j^{-1}(\mathbf{y}-\boldsymbol{\mu}_j)}$$

are multivariate normals $N_d(\boldsymbol{\mu}_j, \boldsymbol{\Omega}_j)$.

- Given data $\mathbf{y}_1, \dots, \mathbf{y}_n$, want to estimate parameters

$$\boldsymbol{\theta} = (\pi_1, \dots, \pi_k, \boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_k, \boldsymbol{\Omega}_1, \dots, \boldsymbol{\Omega}_k).$$

(Incomplete) data log-likelihood is

$$\ln g(\mathbf{y}_1, \dots, \mathbf{y}_n \mid \boldsymbol{\theta}) = \sum_{i=1}^n \ln h(\mathbf{y}_i) = \sum_{i=1}^n \ln \sum_{j=1}^k \pi_j h_j(\mathbf{y}_i \mid \boldsymbol{\mu}_j, \boldsymbol{\Omega}_j).$$

- Let $z_{ij} = I\{\mathbf{y}_i \text{ comes from group } j\}$. Complete data likelihood is

$$f(\mathbf{y}, \mathbf{z}|\boldsymbol{\theta}) = \prod_{i=1}^n \prod_{j=1}^k [\pi_j h_j(\mathbf{y}_i|\boldsymbol{\mu}_j, \boldsymbol{\Omega}_j)]^{z_{ij}}$$

and thus complete log-likelihood is

$$\ln f(\mathbf{y}, \mathbf{z}|\boldsymbol{\theta}) = \sum_{i=1}^n \sum_{j=1}^k z_{ij} [\ln \pi_j + \ln h_j(\mathbf{y}_i|\boldsymbol{\mu}_j, \boldsymbol{\Omega}_j)].$$

- E step: need to evaluate conditional expectation

$$\begin{aligned} & Q(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)}) \\ &= \mathbf{E} \left\{ \sum_{i=1}^n \sum_{j=1}^k z_{ij} [\ln \pi_j + \ln h_j(\mathbf{y}_i|\boldsymbol{\mu}_j, \boldsymbol{\Omega}_j)] \mid \mathbf{Y} = \mathbf{y}, \boldsymbol{\pi}^{(t)}, \boldsymbol{\mu}_1^{(t)}, \dots, \boldsymbol{\mu}_k^{(t)}, \boldsymbol{\Omega}_1^{(t)}, \dots, \boldsymbol{\Omega}_k^{(t)} \right\}. \end{aligned}$$

By Bayes rule, we have

$$\begin{aligned} w_{ij}^{(t)} &:= \mathbf{E}[z_{ij} \mid \mathbf{y}, \boldsymbol{\pi}^{(t)}, \boldsymbol{\mu}_1^{(t)}, \dots, \boldsymbol{\mu}_k^{(t)}, \boldsymbol{\Omega}_1^{(t)}, \dots, \boldsymbol{\Omega}_k^{(t)}] \\ &= \frac{\pi_j^{(t)} h_j(\mathbf{y}_i|\boldsymbol{\mu}_j^{(t)}, \boldsymbol{\Omega}_j^{(t)})}{\sum_{j'=1}^k \pi_{j'}^{(t)} h_{j'}(\mathbf{y}_i|\boldsymbol{\mu}_{j'}^{(t)}, \boldsymbol{\Omega}_{j'}^{(t)})}. \end{aligned}$$

So the Q function becomes

$$\begin{aligned} & Q(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)}) \\ &= \sum_{i=1}^n \sum_{j=1}^k w_{ij}^{(t)} \ln \pi_j + \sum_{i=1}^n \sum_{j=1}^k w_{ij}^{(t)} \left[-\frac{1}{2} \ln \det \boldsymbol{\Omega}_j - \frac{1}{2} (\mathbf{y}_i - \boldsymbol{\mu}_j)^{\top} \boldsymbol{\Omega}_j^{-1} (\mathbf{y}_i - \boldsymbol{\mu}_j) \right]. \end{aligned}$$

- M step: maximizer of the Q function gives the next iterate

$$\begin{aligned} \pi_j^{(t+1)} &= \frac{\sum_i w_{ij}^{(t)}}{n} \\ \boldsymbol{\mu}_j^{(t+1)} &= \frac{\sum_{i=1}^n w_{ij}^{(t)} \mathbf{y}_i}{\sum_{i=1}^n w_{ij}^{(t)}} \\ \boldsymbol{\Omega}_j^{(t+1)} &= \frac{\sum_{i=1}^n w_{ij}^{(t)} (\mathbf{y}_i - \boldsymbol{\mu}_j^{(t+1)}) (\mathbf{y}_i - \boldsymbol{\mu}_j^{(t+1)})^{\top}}{\sum_i w_{ij}^{(t)}}. \end{aligned}$$

See KL Example 11.3.1 for multinomial MLE. See KL Example 11.2.3 for multivariate normal MLE.

- Compare these extremely simple updates to Newton type algorithms!
- Also note the ease of parallel computing with this EM algorithm. See, e.g., Suchard, M. A.; Wang, Q.; Chan, C.; Frelinger, J.; Cron, A. & West, M. Understanding GPU programming for statistical computation: studies in massively parallel massive mixtures. *Journal of Computational and Graphical Statistics*, 2010, 19, 419-438.
- In general, EM/MM algorithms are particularly attractive for parallel computing. See, e.g., H Zhou, K Lange, & M Suchard. (2010) Graphical processing units and high-dimensional optimization, *Statistical Science*, 25:311-324.

21 Lecture 21, Nov 11

Announcements

- HW6 returned. Sketch of solution: <http://hua-zhou.github.io/teaching/st758-2014fall/hw06sol.html>
- HW7 due next Tue Nov 18
- HW8 due Nov 25


Last time

- Nonlinear regression (Gauss-Newton algorithm)
- EM algorithm
- Example: finite mixture model


Today

- MM algorithm

MM algorithm (KL Ch12)



- Recall our picture for understanding the ascent property of EM



- EM as a minorization-maximization (MM) algorithm
 - The Q function constitutes a *minorizing* function of the objective function up to an additive constant

$$L(\boldsymbol{\theta}) \geq Q(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)}) + c^{(t)} \quad \text{for all } \boldsymbol{\theta}$$

$$L(\boldsymbol{\theta}^{(t)}) = Q(\boldsymbol{\theta}^{(t)}|\boldsymbol{\theta}^{(t)}) + c^{(t)}$$

- *Maximizing* the Q function generates an ascent iterate $\boldsymbol{\theta}^{(t+1)}$
- Questions:
 - Is EM principle only limited to maximizing likelihood model?
 - Is there any other way to produce such surrogate function?
 - Can we flip the picture and apply same principle to *minimization* problem?

These thoughts lead to a powerful tool – MM principle (Lange et al., 2000).
 Lange, K., Hunter, D. R., and Yang, I. (2000). Optimization transfer using surrogate objective functions. *J. Comput. Graph. Statist.*, 9(1):159. With discussion, and a rejoinder by Hunter and Lange.

- For maximization of $f(\boldsymbol{\theta})$ – minorization-maximization (MM) algorithm
 - **Minorization step:** Construct a surrogate function $g(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)})$ such that

$$f(\boldsymbol{\theta}) \geq g(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)}) \quad (\text{dominance condition})$$

$$f(\boldsymbol{\theta}^{(t)}) = g(\boldsymbol{\theta}^{(t)}|\boldsymbol{\theta}^{(t)}) \quad (\text{tangent condition}).$$

– Maximization step:

$$\boldsymbol{\theta}^{(t+1)} = \operatorname{argmax} g(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)}).$$

- *Ascent* property of minorization-maximization algorithm

$$f(\boldsymbol{\theta}^{(t)}) = g(\boldsymbol{\theta}^{(t)}|\boldsymbol{\theta}^{(t)}) \leq g(\boldsymbol{\theta}^{(t+1)}|\boldsymbol{\theta}^{(t)}) \leq f(\boldsymbol{\theta}^{(t+1)}).$$

- EM is a special case of minorization-maximization (MM) algorithm.


- For minimization of $f(\boldsymbol{\theta})$ – majorization-minimization (MM) algorithm

– Majorization step: Construct a surrogate function $g(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)})$ such that

$$\begin{aligned} f(\boldsymbol{\theta}) &\leq g(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)}) \quad (\text{dominance condition}) \\ f(\boldsymbol{\theta}^{(t)}) &= g(\boldsymbol{\theta}^{(t)}|\boldsymbol{\theta}^{(t)}) \quad (\text{tangent condition}). \end{aligned}$$

– Minimization step:


$$\boldsymbol{\theta}^{(t+1)} = \operatorname{argmin} g(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)}).$$



- Similarly we have the *descent* property of majorization-minimization algorithm.
- Same convergence theory as EM.

- Rational of the MM principle for developing optimization algorithms
 - Free the derivation from missing data structure.
 - Avoid matrix inversion.
 - Linearize an optimization problem.
 - Deal gracefully with certain equality and inequality constraints.
 - Turn a non-differentiable problem into a smooth problem.
 - Separate the parameters of a problem (perfect for massive, *fine-scale* parallelization).
- Generic methods for majorization and minorization – *inequalities*
 - Jensen’s (information) inequality – EM algorithms
 - The Cauchy-Schwartz inequality - multidimensional scaling
 - Supporting hyperplane property of a convex function
 - Arithmetic-geometric mean inequality
 - Quadratic upper bound principle - Böhning and Lindsay
 - ...
- Numerous examples in KL chapter 12.
- History: the name *MM algorithm* originates from the discussion (by Xiaoli Meng) and rejoinder of the Lange et al. (2000) paper.

Example: PET imaging



- Data: tube readings $\mathbf{y} = (y_1, \dots, y_d)$.
- Estimate: photon emission intensities (pixels) $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_p)$.

- Poisson Model:

$$Y_i \sim \text{Poisson} \left(\sum_{j=1}^p c_{ij} \lambda_j \right),$$

where c_{ij} is the (pre-calculated) cond. prob. that a photon emitted by j -th pixel is detected by i -th tube.

- Log-likelihood:

$$L(\boldsymbol{\lambda}|\mathbf{y}) = \sum_i \left[y_i \ln \left(\sum_j c_{ij} \lambda_j \right) - \sum_j c_{ij} \lambda_j \right] + \text{const.}$$

Essentially a Poisson regression with constraint $\lambda_j \geq 0$.

- Which algorithm?

– Fisher scoring (IRWLS) = Newton.

Need to solve a large linear system at each iteration ☹


– EM algorithm: Lange and Carson (1984), Vardi et al. (1985)

$$\lambda_j^{(t+1)} = \lambda_j^{(t)} \sum_i \frac{y_i c_{ij}}{\sum_k c_{ik} \lambda_k^{(t)}}.$$

Scales well with data size. Converges to the global maximum under mild conditions.

– Exercise: derive the EM algorithm. (Hint: missing data $z_{ij} = \#$ of photons emitted from pixel i and received by detector j .)

- Issues: *grainy image* and *slow convergence*



- Regularized log-likelihood for smoother image:

$$\begin{aligned}
L(\boldsymbol{\lambda}|\mathbf{y}) &- \frac{\mu}{2} \sum_{\{j,k\} \in \mathcal{N}} (\lambda_j - \lambda_k)^2 \\
&= \sum_i \left[y_i \ln \left(\sum_j c_{ij} \lambda_j \right) - \sum_j c_{ij} \lambda_j \right] - \frac{\mu}{2} \sum_{\{j,k\} \in \mathcal{N}} (\lambda_j - \lambda_k)^2,
\end{aligned}$$

where $\mu \geq 0$ is a tuning constant.

- EM algorithm does not (or is hard to) apply to the regularization term. Let's derive an MM algorithm.
- Minorization step:

– By concavity of the $\ln s$ function

$$\begin{aligned}
\ln \left(\sum_j c_{ij} \lambda_j \right) &= \ln \left(\sum_j \frac{c_{ij} \lambda_j^{(t)}}{\sum_{j'} c_{ij'} \lambda_{j'}^{(t)}} \cdot \frac{\sum_{j'} c_{ij'} \lambda_{j'}^{(t)}}{c_{ij} \lambda_j^{(t)}} \cdot c_{ij} \lambda_j \right) \\
&\geq \sum_j \frac{c_{ij} \lambda_j^{(t)}}{\sum_{j'} c_{ij'} \lambda_{j'}^{(t)}} \ln \left(\frac{\sum_{j'} c_{ij'} \lambda_{j'}^{(t)}}{c_{ij} \lambda_j^{(t)}} c_{ij} \lambda_j \right) \\
&= \sum_j \frac{c_{ij} \lambda_j^{(t)}}{\sum_{j'} c_{ij'} \lambda_{j'}^{(t)}} \ln \lambda_j + c^{(t)}.
\end{aligned}$$

Remark: this minorization depends on the positivity of both c_{ij} and λ_j .

– By concavity of the $-s^2$ function

$$\begin{aligned}
-(\lambda_j - \lambda_k)^2 &= - \left(\frac{2\lambda_j - \lambda_j^{(t)} - \lambda_k^{(t)}}{2} + \frac{-2\lambda_k + \lambda_j^{(t)} + \lambda_k^{(t)}}{2} \right)^2 \\
&\geq -\frac{1}{2} (2\lambda_j - \lambda_j^{(t)} - \lambda_k^{(t)})^2 - \frac{1}{2} (2\lambda_k - \lambda_j^{(t)} - \lambda_k^{(t)})^2.
\end{aligned}$$

– Combining minorizing terms gives an overall surrogate function

$$\begin{aligned}
g(\boldsymbol{\lambda}|\boldsymbol{\lambda}^{(t)}) &= \sum_i y_i \sum_j \frac{c_{ij} \lambda_j^{(t)}}{\sum_{j'} c_{ij'} \lambda_{j'}^{(t)}} \ln \lambda_j - \sum_i \sum_j c_{ij} \lambda_j \\
&\quad - \frac{\mu}{4} \sum_{\{j,k\} \in \mathcal{N}} [(2\lambda_j - \lambda_j^{(t)} - \lambda_k^{(t)})^2 + (2\lambda_k - \lambda_j^{(t)} - \lambda_k^{(t)})^2].
\end{aligned}$$

- Maximization step:

- $g(\boldsymbol{\lambda}|\boldsymbol{\lambda}^{(t)})$ is trivial to maximize because all λ_j are separated!
- Solving for the root of

$$\begin{aligned} & \frac{\partial}{\partial \lambda_j} g(\boldsymbol{\lambda}|\boldsymbol{\lambda}^{(t)}) \\ &= \left(\sum_i y_i \frac{c_{ij} \lambda_j^{(t)}}{\sum_{j'} c_{ij} \lambda_{j'}^{(t)}} \right) \lambda_j^{-1} - \sum_i c_{ij} - \mu \sum_{k \in \mathcal{N}_j} (2\lambda_j - \lambda_j^{(t)} - \lambda_k^{(t)}) \\ &= 0 \end{aligned}$$

gives $\lambda_j^{(t+1)}$.

- MM algorithm for PET:

Initialize: $\lambda_j^{(0)} = 1$

repeat

$$z_{ij}^{(t)} = (y_i c_{ij} \lambda_j^{(t)}) / (\sum_k c_{ik} \lambda_k^{(t)})$$

for $j = 1$ to p **do**

$$a = -2\mu|\mathcal{N}_j|, b = \mu(|\mathcal{N}_j|\lambda_j^{(t)} + \sum_{k \in \mathcal{N}_j} \lambda_k^{(t)}) - 1, c = \sum_i z_{ij}^{(t)}$$

$$\lambda_j^{(t+1)} = (-b - \sqrt{b^2 - 4ac}) / (2a)$$

end for

until convergence occurs

- Parameter constraints $\lambda_j \geq 0$ are satisfied when start from positive initial values.
- The loop for updating pixels can be carried out independently – *massive parallelism*.
- A simulation example with $n = 2016$ and $p = 4096$ (provided by Ravi Varadhan)
CPU: i7 @ 3.20GHZ (1 thread), implemented using BLAS in the GNU Scientific Library (GSL)
GPU: NVIDIA GeForce GTX 580, implemented using cuBLAS

Penalty μ	CPU			GPU			
	Iters	Time	Function	Iters	Time	Function	Speedup
0	100000	11250	-7337.152765	100000	140	-7337.153387	80
10^{-7}	24506	2573	-8500.082605	24506	35	-8508.112249	74
10^{-6}	6294	710	-15432.45496	6294	9	-15432.45586	79
10^{-5}	589	67	-55767.32966	589	0.8	-55767.32970	84

22 Lecture 22, Nov 13

Announcements

- HW7 due next Tue Nov 18
- HW8 due Nov 25
- HW9 (simulation project) will be posted this week. Due Dec 9, 2014 @ 11A.

Last time

- MM algorithm
- Example: PET imaging

Today

- HW6 review
- Principles of Monte carlo simulation studies
- EM/MM example: Netflix problem

Feedback on HW6

- Who cares?

TABLE 3.4. Allele Counts in Four Subpopulations

Allele	White	Black	Chicano	Asian
5	2	0	0	0
6	84	50	80	16
7	59	137	128	40
8	41	78	26	8
9	53	54	55	68
10	131	51	95	14
11	2	0	0	7
12	0	0	0	1
Total $2n$	372	370	384	154

TABLE 3.5. Classical and Bayesian Allele Frequency Estimates

Allele	White	Black	Chicano	Asian
5	.0054 .0053	.0000 .0003	.0000 .0003	.0000 .0006
6	.2258 .2227	.1351 .1380	.2083 .2064	.1039 .1147
7	.1586 .1667	.3703 .3645	.3333 .3301	.2597 .2630
8	.1102 .1105	.2108 .2045	.0677 .0707	.0519 .0609
9	.1425 .1465	.1459 .1498	.1432 .1471	.4416 .4073
10	.3522 .3424	.1378 .1421	.2474 .2445	.0909 .1070
11	.0054 .0057	.0000 .0007	.0000 .0007	.0455 .0404
12	.0000 .0002	.0000 .0002	.0000 .0002	.0065 .0061
Sample Size $2n$	372	370	384	154

- Bayesian approach to estimate a multinomial parameter \mathbf{p}
 - Data likelihood: $\mathbf{x}|\mathbf{p} \sim \text{multinomial}(\mathbf{p})$
 - Prior: $\mathbf{p} \sim \text{Dirichlet}(\boldsymbol{\alpha})$
 - Posterior: $\mathbf{p}|\mathbf{x} \sim \text{Dirichlet}(\boldsymbol{\alpha} + \mathbf{x})$
 - Do estimation and inference of \mathbf{p} based on the posterior distribution


But what value of $\boldsymbol{\alpha}$ to use in the prior?

- Empirical Bayes idea:
 - Estimate $\boldsymbol{\alpha}$ from data, say by maximizing the marginal likelihood of

$$\mathbf{x}_i \sim \text{DirMult}(\boldsymbol{\alpha}).$$

- Then estimate \mathbf{p} by posterior mean $\hat{\mathbf{p}}_{\text{EB}} = (\mathbf{x} + \hat{\boldsymbol{\alpha}})/(|\mathbf{x}| + |\hat{\boldsymbol{\alpha}}|)$.

For a celebrated binomial estimation problem (batting averages of major league baseball players), see Efron and Morris (1973, 1977).



- The Bayes estimate (under certain conditions) enjoys both good asymptotic and finite sample properties.

“borrow information across populations”, “shrinkage”, “learning from the experience of the others”

The Annals of Statistics
1986, Vol. 14, No. 1, 1–26

SPECIAL INVITED PAPER

ON THE CONSISTENCY OF BAYES ESTIMATES

BY PERSI DIACONIS¹ AND DAVID FREEDMAN²

Stanford University and University of California, Berkeley

We discuss frequency properties of Bayes rules, paying special attention to consistency. Some new and fairly natural counterexamples are given, involving nonparametric estimates of location. Even the Dirichlet prior can lead to inconsistent estimates if used too aggressively. Finally, we discuss reasons for Bayesians to be interested in frequency properties of Bayes rules. As a part of the discussion we give a subjective equivalent to consistency and compute the derivative of the map taking priors to posteriors.

1. Consistency of Bayes rules. One of the basic problems in statistics can be put this way. Data is collected following a probability model with unknown parameters; the parameters are to be estimated from the data. Often, there is prior information about the parameters, for example, their probable sign or order of magnitude. Many statisticians express such information in the form of a prior probability over the unknown parameters. Estimates based on prior probabilities will be called Bayes estimates in what follows.

- Machine learning applications: Handwritten digit recognition, text mining, email spam detection, ...




FIGURE 13.9. Examples of grayscale images of handwritten digits.

TABLE 7
Confusion Matrix for the WebKB4 Data Set Using MDD Mixture

	Course	Faculty	Project	Student
Course	826	47	33	24
Faculty	44	901	38	141
Project	31	47	374	52
Student	41	52	43	1505

- In HW6/HW7, we estimate α by MLE.
- HW6 (Newton’s method)

- Check solution. It's not optimal but you may learn things. <http://hua-zhou.github.io/teaching/st758-2014fall/hw06sol.html>
- Implementation of density function: deal with $\ln \Gamma(0) - \ln \Gamma(0) = \infty - \infty = NaN$ issue. Is the first formulation more efficient for computation (log instead of log gamma evaluation)?
- Implementation of MLE function. Allowing observation weights is always a good idea.
- Implementation of gradient/score function. Is the first formulation more efficient (free of digamma evaluations)?
- Special structure in the observed information matrix

$$-d^2 L(\boldsymbol{\alpha}) = \mathbf{D} - c\mathbf{1}\mathbf{1}^T,$$

where $\mathbf{D} = \text{diag}(d_1, \dots, d_d)$ and

$$d_j = \sum_{i=1}^n \sum_{k=0}^{x_{ij}-1} \frac{1}{(\alpha_j + k)^2}, \quad j = 1, \dots, d,$$

$$c = \sum_{i=1}^n \sum_{k=0}^{|\mathbf{x}_i|} \frac{1}{(|\boldsymbol{\alpha}| + k)^2}.$$

Is the first formulation more efficient (free of trigamma evaluations)?

- By Sherman-Morrison

$$[-d^2 L(\boldsymbol{\alpha})]^{-1} = \mathbf{D}^{-1} + \frac{1}{c^{-1} - \sum_i d_i^{-1}} (\mathbf{D}^{-1}\mathbf{1})(\mathbf{1}^T \mathbf{D}^{-1}),$$

which is pd if and only if $c^{-1} > \sum_i d_i^{-1}$.

- Ad hoc approximation of Hessian: use $\tilde{c} = 0.95(\sum_i d_i^{-1})^{-1}$ if it's not satisfied.
- Newton's direction is computed fast: $O(d)$ flops. No need to store any matrix! Why do I still see `solve(D)` or matrix-vector multiplication?
- How to calculate Fisher information matrix?
- Deal with boundary constraint. Choose initial step size s such that all components of

$$\boldsymbol{\alpha}^{(t)} + s\Delta\boldsymbol{\alpha}$$

are positive.

- What’s that damn MoM starting point? If $\mathbf{P} = (P_1, \dots, P_d)$ is from Dirichlet with parameter $\boldsymbol{\alpha}$, then

$$\mathbf{E}(P_j) = \frac{\alpha_j}{|\boldsymbol{\alpha}|}, \quad \mathbf{E}(P_j^2) = \frac{\alpha_j(\alpha_j + 1)}{|\boldsymbol{\alpha}|(|\boldsymbol{\alpha}| + 1)}.$$

Therefore

$$\sum_j \frac{\mathbf{E}(P_j^2)}{\mathbf{E}(P_j)} = \frac{|\boldsymbol{\alpha}| + d}{|\boldsymbol{\alpha}| + 1}.$$

We estimate the left hand side by

$$\hat{\rho} := \sum_j \frac{\sum_i (x_{ij}/|\mathbf{x}_i|)^2}{\sum_i (x_{ij}/|\mathbf{x}_i|)}$$

and then solve for $|\boldsymbol{\alpha}^{(0)}| = (d - \hat{\rho})/(\hat{\rho} - 1)$. Then we initialize

$$\alpha_j^{(0)} = |\boldsymbol{\alpha}^{(0)}| \frac{\sum_i x_{ij}}{\sum_i |\mathbf{x}_i|} = \left(\frac{d - \hat{\rho}}{\hat{\rho} - 1} \right) \left(\frac{\sum_i x_{ij}}{\sum_i |\mathbf{x}_i|} \right).$$

Is there any proof $(d - \hat{\rho})/(\hat{\rho} - 1)$ is always positive? Anyway we just need a good starting point. Other heuristics should also work.

- HW7: EM/MM algorithm for maximizing

$$L(\boldsymbol{\alpha}) = \sum_{i=1}^n \ln \binom{|\mathbf{x}_i|}{\mathbf{x}_i} + \sum_{i=1}^n \sum_{j=1}^d \sum_{k=0}^{x_{ij}-1} \ln(\alpha_j + k) - \sum_{i=1}^n \sum_{k=0}^{|\mathbf{x}_i|-1} \ln(|\boldsymbol{\alpha}| + k)$$

Hint: Apply Jensen’s inequality to $\ln(\alpha_j + k)$ and supporting hyperplane inequality to $\ln(|\boldsymbol{\alpha}| + k)$.

- Examining Newton iterates, we are delighted to see its fast (quadratic) convergence. In HW7, you might see another scenario (linear convergence) for MM iterates. On the other hand, for machine learning problem, do we really need a super-accurate solution?
- Main messages:
 - basics of Newton’s method (pd approximation, line search)
 - examine and exploit problem structure whenever possible

- EM and MM turn out to be different
- Comparing Newton versus MM: number of iterations, human efforts, efficiency, ...

(Optimization is a) fascinating blend of theory and computation, heuristics and rigor.

Roger Fletcher

There is simply no such thing as a universal ‘gold standard’ when it comes to algorithms.

Unknown Reviewer

- Implementation details: `lgamma`, `digamma`, `trigamma`, `psigamma` in R are vectorized function, vectorize code (numerical linear algebra is preferred over `apply()`), ...

Some black-box optimization routines in R

General-purpose Optimization

Description

General-purpose optimization based on Nelder–Mead, quasi-Newton and conjugate-gradient algorithms. It includes an option for box-constrained optimization and simulated annealing.

Usage

```
optim(par, fn, gr = NULL, ...,
      method = c("Nelder-Mead", "BFGS", "CG", "L-BFGS-B", "SANN",
                 "Brent"),
      lower = -Inf, upper = Inf,
      control = list(), hessian = FALSE)
```

Optimization using PORT routines

Description

Unconstrained and box-constrained optimization using PORT routines.

Usage

```
nlminb(start, objective, gradient = NULL, hessian = NULL, ...,  
        scale = 1, control = list(), lower = -Inf, upper = Inf)
```

23 Lecture 23, Nov 18

Announcements

- HW7 (EM/MM for Dirichlet-multinomial) due today: submit hardcopy + email code (`LastFirstHW7.R` or `LastFirstHW7.Rmd`)
- HW8 (ranking MLB teams) due next Tue Nov 25
- HW9 (Monte Carlo simulation project, 100 pts) posted. Due Dec 9 @ 11A. Send both your report (pdf) and code (`R` or `Rmd`) by email.
 - Read posted course materials
 - Start early and send me your questions or draft by Nov 30
 - Discussion on Dec 2 (last lecture)

Last time

- HW6 review (Newton's method in action)

Today

- Last EM/MM example: Netflix problem

Example: Netflix and matrix completion

- Snapshot of the kind of data collected by Netflix. Only 100,480,507 ratings (1.2% entries of the 480K-by-18K matrix) are observed

ID	movie 1	movie 2	movie 3	movie 4	movie 5	movie 6	...	movie 17,770
user 1	5	3	4	3	3	NA	...	1
user 2	4	NA	NA	NA	NA	NA	...	NA
user 3	NA	NA	NA	NA	NA	NA	...	NA
user 4	4	NA	NA	NA	NA	2	...	4
user 5	NA	NA	NA	5	NA	NA	...	NA
user 6	3	NA	NA	5	1	NA	...	3
user 7	NA	NA	NA	NA	NA	NA	...	NA
user 8	5	NA	5	NA	NA	NA	...	NA
user 9	NA	NA	NA	NA	3	NA	...	NA
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
user 480,189	NA	5	NA	NA	NA	NA	...	NA

- Netflix challenge: impute the unobserved ratings for personalized recommendation. http://en.wikipedia.org/wiki/Netflix_Prize



- *Matrix completion problem.* Observe a very sparse matrix $\mathbf{Y} = (y_{ij})$. Want to impute all the missing entries. It is possible only when the matrix is structured, e.g., of *low rank*.
- Let $\Omega = \{(i, j) : \text{observed entries}\}$ index the observed entries and $P_{\Omega}(\mathbf{M})$ denote the projection of matrix \mathbf{M} to Ω . The problem

$$\min_{\text{rank}(\mathbf{X}) \leq r} \frac{1}{2} \|P_{\Omega}(\mathbf{Y}) - P_{\Omega}(\mathbf{X})\|_F^2 = \frac{1}{2} \sum_{(i,j) \in \Omega} (y_{ij} - x_{ij})^2$$

unfortunately is non-convex and difficult.

- *Convex relaxation* (Mazumder et al., 2010)

$$\min_{\mathbf{X}} f(\mathbf{X}) = \frac{1}{2} \|P_{\Omega}(\mathbf{Y}) - P_{\Omega}(\mathbf{X})\|_F^2 + \lambda \|\mathbf{X}\|_*,$$

where $\|\mathbf{X}\|_* = \|\sigma(\mathbf{X})\|_1 = \sum_i \sigma_i(\mathbf{X})$ is the nuclear norm.

- Majorization step:

$$\begin{aligned}
f(\mathbf{X}) &= \frac{1}{2} \sum_{(i,j) \in \Omega} (y_{ij} - x_{ij})^2 + \frac{1}{2} \sum_{(i,j) \notin \Omega} 0 + \lambda \|\mathbf{X}\|_* \\
&\leq \frac{1}{2} \sum_{(i,j) \in \Omega} (y_{ij} - x_{ij})^2 + \frac{1}{2} \sum_{(i,j) \notin \Omega} (x_{ij}^{(t)} - x_{ij})^2 + \lambda \|\mathbf{X}\|_* \\
&= \frac{1}{2} \|\mathbf{X} - \mathbf{Z}^{(t)}\|_F^2 + \lambda \|\mathbf{X}\|_* \\
&= g(\mathbf{X}|\mathbf{X}^{(t)}),
\end{aligned}$$

where $\mathbf{Z}^{(t)} = P_\Omega(\mathbf{Y}) + P_{\Omega^\perp}(\mathbf{X}^{(t)})$. “fill in missing entries by current imputation”

- Minimization step:

Rewrite the surrogate function

$$g(\mathbf{X}|\mathbf{X}^{(t)}) = \frac{1}{2} \|\mathbf{X}\|_F^2 - \text{tr}(\mathbf{X}^T \mathbf{Z}^{(t)}) + \frac{1}{2} \|\mathbf{Z}^{(t)}\|_F^2 + \lambda \|\mathbf{X}\|_*.$$

Let σ_i be the singular values of \mathbf{X} and ω_i be the singular values of $\mathbf{Z}^{(t)}$. Observe

$$\begin{aligned}
\|\mathbf{X}\|_F^2 &= \|\sigma(\mathbf{X})\|_2^2 = \sum_i \sigma_i^2 \\
\|\mathbf{Z}^{(t)}\|_F^2 &= \|\sigma(\mathbf{Z}^{(t)})\|_2^2 = \sum_i \omega_i^2
\end{aligned}$$

and by the Fan-von Neuman’s inequality

$$\text{tr}(\mathbf{X}^T \mathbf{Z}^{(t)}) \leq \sum_i \sigma_i \omega_i$$

with equality achieved if and only if the left and right singular vectors of the two matrices coincide. Thus we can choose \mathbf{X} to have same singular vectors as $\mathbf{Z}^{(t)}$ and

$$\begin{aligned}
g(\mathbf{X}|\mathbf{X}^{(t)}) &= \frac{1}{2} \sum_i \sigma_i^2 - \sum_i \sigma_i \omega_i + \frac{1}{2} \omega_i^2 + \lambda \sum_i \sigma_i \\
&= \frac{1}{2} \sum_i (\sigma_i - \omega_i)^2 + \lambda \sum_i \sigma_i,
\end{aligned}$$

with minimizer given by $\sigma_i^{(t+1)} = (\omega_i - \lambda)_+$. “Singular value thresholding”

- Algorithm for matrix completion:

Initialize $\mathbf{X}^{(0)} \in \mathbb{R}^{m \times n}$

repeat

$$\mathbf{Z}^{(t+1)} \leftarrow P_{\Omega}(\mathbf{Y}) + P_{\Omega^{\perp}}(\mathbf{X}^{(t)})$$

$$\text{Compute SVD: } \mathbf{U} \text{diag}(\mathbf{w}) \mathbf{V}^{\top} \leftarrow \mathbf{Z}^{(t+1)}$$

$$\mathbf{X}^{(t+1)} \leftarrow \mathbf{U} \text{diag}[(\mathbf{w} - \lambda)_+] \mathbf{V}^{\top}$$

until objective value converges

- “Golub-Kahan-Reinsch” algorithm takes $4m^2n + 8mn^2 + 9n^3$ flops for a $m \geq n$ matrix and is not going to work for 480K-by-18K Netflix matrix.

Notice only top singular values are needed since low rank solutions are achieved by large λ . Lanczos/Arnoldi algorithm is the way to go. Matrix-vector multiplication $\mathbf{Z}^{(t)}\mathbf{v}$ is fast (why?)

24 Lecture 24, Nov 20

Announcements

- HW7 (50pts) returned. Sketch of solution: <http://hua-zhou.github.io/teaching/st758-2014fall/hw07sol.html> Feedback:
 - Vectorize code by `outer` + `apply` or `table` + `apply`
 - Pre-compute certain quantities s_{jk} and r_k
 - Convergence criterion: relative change in objective values, or first order optimality condition
- HW8 (ranking MLB teams) due next Tue Nov 25.
- HW9 (simulation project) due Dec 9 @ 11A.

Last time

- EM/MM example: Netflix problem

Today

- Quasi-Newton method.
- Conjugate gradient method.

Quasi-Newton methods (KL 14.9)

- Quasi-Newton is probably the most successful “black-box” optimizers in use. E.g., implemented in the general purpose optimization routine `optim()` in R.
- Consider the general Newton scheme for minimizing $f(\mathbf{x})$, $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^p$:

$$\mathbf{x}^{(t+1)} = \mathbf{x}^{(t)} - s[\mathbf{A}^{(t)}]^{-1}\nabla f(\mathbf{x}^{(t)}),$$

where $\mathbf{A}^{(t)}$ a pd approximation of the Hessian $d^2 f(\mathbf{x}^{(t)})$.

- Pros: fast (quadratic) convergence

- Cons: compute and store Hessian at each iteration (usually $O(np^2)$ cost in statistical problems), solving a linear system ($O(p^3)$ cost in general), human efforts (derive and implement gradient and Hessian, pd approximation, ...)
- Any pd \mathbf{A} gives a descent algorithm – tradeoff between convergence rate and cost per iteration.
- Setting $\mathbf{A} = \mathbf{I}$ leads to the *steepest descent* algorithm. Picture: slow convergence (zig-zagging) of steepest descent (with exact line search) for minimizing a convex quadratic function (ellipse).




Figure 10.6.1. (a) Steepest descent method in a long, narrow “valley.” While more efficient than the strategy of Figure 10.5.1, steepest descent is nonetheless an inefficient strategy, taking many steps to reach the valley floor. (b) Magnified view of one step: A step starts off in the local gradient direction, perpendicular to the contour lines, and traverses a straight line until a local minimum is reached, where the traverse is parallel to the local contour lines.

How many steps does the Newton’s method take for a convex quadratic f ?

- Idea of Quasi-Newton: No analytical Hessian is required (still need gradient). Update approximate Hessian \mathbf{A} according to the *secant condition*

$$\nabla f(\mathbf{x}^{(t-1)}) - \nabla f(\mathbf{x}^{(t)}) \approx d^2 f(\mathbf{x}^{(t)})(\mathbf{x}^{(t-1)} - \mathbf{x}^{(t)}).$$

Instead of computing \mathbf{A} from scratch at each iteration, we update an approximation \mathbf{A} to $d^2 f(\mathbf{x})$ which satisfies (1) p.d., (2) the secant condition, and (3) closest to the previous approximation.

- Super-linear convergence, compared to the quadratic convergence of Newton’s method. But each iteration only takes $O(p^2)$!
- History: Davidon was a physicist at Argonne National Lab in 50s and proposed the very first idea of quasi-Newton method in 1959. Later studied, implemented, and popularized by Fletcher and Powell. Davidon’s original paper was not accepted for publication ☹; 30 years later it appeared as the first article in the first issue of *SIAM Journal of Optimization* (Davidon, 1991).

William C. Davidon

From Wikipedia, the free encyclopedia

William Cooper Davidon (1927–November 8, 2013) was an American professor of physics and mathematics, and peace activist. He was the mastermind of the March 8, 1971 F.B.I. office breakin, in Media, Pennsylvania, and the informal leader of the Citizens' Commission to Investigate the FBI.

Contents

- 1 Life
- 2 Activism
- 3 Family
- 4 References
- 5 External links


Life

Davidon was born in Fort Lauderdale, Florida in 1927. He attended Purdue University, and graduated from the University of Chicago with a Ph.D. in 1957.^[1]


From 1954 to 1956, he was a research associate at the Enrico Fermi Institute. From 1956 to 1961, he was an associate physicist at the Argonne National Laboratory. He was professor of physics at Haverford College, beginning in 1961, and then Professor of Mathematics. He retired in 1991. He was a 1966 Fulbright Scholar.^[2]

Davidon moved to Highlands Ranch, Colorado, in 2010. He died November 8, 2013, of Parkinson's disease.

William C. Davidon



Born	1927 Fort Lauderdale
Died	November 8, 2013 Highlands Ranch, Colorado
Nationality	American
Occupation	physics professor
Known for	Citizens' Commission to Investigate the FBI



- Davidon-Fletcher-Powell (DFP) rank-2 update. Solve

$$\begin{aligned} & \text{minimize} && \| \mathbf{A} - \mathbf{A}^{(t)} \|_F \\ & \text{subject to} && \mathbf{A} = \mathbf{A}^T \\ & && \mathbf{A}(\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}) = \nabla f(\mathbf{x}^{(t)}) - \nabla f(\mathbf{x}^{(t-1)}) \end{aligned}$$

for the next approximation $\mathbf{A}^{(t+1)}$. The solution is a low rank (rank 1 or rank 2) update of $\mathbf{A}^{(t)}$. The inverse is too thanks to Sherman-Morrison-Woodbury! $O(p^2)$ operations. Need to store a p -by- p dense matrix. Positive definiteness is guaranteed by the same trick you used in HW6! See KL 14.9 for details.

- Broyden-Fletcher-Goldfarb-Shanno (BFGS) rank 2 update is considered by many the most effective among all quasi-Newton updates. BFGS imposes secant condition on the inverse of Hessian $\mathbf{H} = \mathbf{A}^{-1}$.

$$\begin{aligned} & \text{minimize} && \|\mathbf{H} - \mathbf{H}^{(t)}\|_F \\ & \text{subject to} && \mathbf{H} = \mathbf{H}^\top \\ & && \mathbf{H}[\nabla f(\mathbf{x}^{(t)}) - \nabla f(\mathbf{x}^{(t-1)})] = \mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}. \end{aligned}$$

Again $\mathbf{H}^{(t+1)}$ is a rank two update of $\mathbf{H}^{(t)}$. $O(p^2)$ operations. Still need to store a dense p -by- p matrix.

- Limited-memory BFGS (L-BFGS). Only store the secant pairs. Particularly useful for large scale optimization.
- L-BFGS-B: with box-constraints. Implemented in the general purpose optimization routine `optim()` in R.
- How to set the initial approximation $\mathbf{A}^{(0)}$? Identity or Hessian (if pd) or Fisher information matrix at starting point.

(Linear) Conjugate gradient method

- (Linear) Conjugate gradient is the top-notch iterative method for solving large, structured linear systems $\mathbf{Ax} = \mathbf{b}$. Earlier we talked about Jacobi and Gauss-Seidel as the more classical iterative solvers.

Table 1. Kershaw's results for a fusion problem.

Method	Number of iterations
Gauss Seidel	208,000
Block successive overrelaxation methods	765
Incomplete Cholesky conjugate gradients	25

- *Linear* conjugate gradient method: for solving large linear systems of equations. History: Hestenes and Stiefel in 50s.

- Solve linear equation $\mathbf{Ax} = \mathbf{b}$, where $\mathbf{A} \in \mathbb{R}^{n \times n}$ is pd, is equivalent to

$$\text{minimize } f(\mathbf{x}) = \frac{1}{2} \mathbf{x}^\top \mathbf{Ax} - \mathbf{b}^\top \mathbf{x}.$$

Note $\nabla f(\mathbf{x}) = \mathbf{Ax} - \mathbf{b} =: \mathbf{r}(\mathbf{x})$.

- Consider a simple idea: coordinate descent, that is to update x_j alternately. Same as the Gauss-Seidel iteration.




Figure 9.1
Coordinate search method makes slow progress on this function of two variables.

- A set of vectors $\{\mathbf{p}^{(0)}, \dots, \mathbf{p}^{(l)}\}$ is said to be conjugate wrt \mathbf{A} if

$$\mathbf{p}_i^\top \mathbf{A} \mathbf{p}_j = 0, \quad \text{for all } i \neq j.$$

- *Conjugate direction* method: Given a set of conjugate vectors $\{\mathbf{p}^{(0)}, \dots, \mathbf{p}^{(l)}\}$, at iteration t , we search along the conjugate direction $\mathbf{p}^{(t)}$

$$\mathbf{x}^{(t+1)} = \mathbf{x}^{(t)} + \alpha^{(t)} \mathbf{p}^{(t)},$$

where

$$\alpha^{(t)} = -\frac{\mathbf{r}^{(t) \top} \mathbf{p}^{(t)}}{\mathbf{p}^{(t) \top} \mathbf{A} \mathbf{p}^{(t)}}.$$

- Theorem: $\mathbf{x}^{(t)}$ converges to the solution in at most n steps.

Intuition: Look at graph.




Figure 5.1 Successive minimizations along the coordinate directions find the minimizer of a quadratic with a diagonal Hessian in n iterations.

- *Conjugate gradient* method. Idea: generate $\mathbf{p}^{(t)}$ using only $\mathbf{p}^{(t-1)}$

$$\mathbf{p}^{(t)} = -\mathbf{r}^{(t)} + \beta^{(t)}\mathbf{p}^{(t-1)},$$

where $\beta^{(t)}$ is determined by the conjugacy condition $\mathbf{p}^{(t-1)\top} \mathbf{A} \mathbf{p}^{(t)} = 0$

$$\beta^{(t)} = \frac{\mathbf{r}^{(t)\top} \mathbf{A} \mathbf{p}^{(t-1)}}{\mathbf{p}^{(t-1)\top} \mathbf{A} \mathbf{p}^{(t-1)}}.$$

- CG algorithm (preliminary version):

Given $\mathbf{x}^{(0)}$

Initialize: $\mathbf{r}^{(0)} \leftarrow \mathbf{A} \mathbf{x}^{(0)} - \mathbf{b}$, $\mathbf{p}^{(0)} \leftarrow -\mathbf{r}^{(0)}$, $t = 0$

while $\mathbf{r}^{(0)} \neq \mathbf{0}$ **do**

$$\alpha^{(t)} \leftarrow -\frac{\mathbf{r}^{(t)\top} \mathbf{p}^{(t)}}{\mathbf{p}^{(t)\top} \mathbf{A} \mathbf{p}^{(t)}}$$

$$\mathbf{x}^{(t+1)} \leftarrow \mathbf{x}^{(t)} + \alpha^{(t)} \mathbf{p}^{(t)}$$

$$\mathbf{r}^{(t+1)} \leftarrow \mathbf{A} \mathbf{x}^{(t+1)} - \mathbf{b}$$

$$\beta^{(t+1)} \leftarrow \frac{\mathbf{r}^{(t+1)\top} \mathbf{A} \mathbf{p}^{(t)}}{\mathbf{p}^{(t)\top} \mathbf{A} \mathbf{p}^{(t)}}$$

$$\mathbf{p}^{(t+1)} \leftarrow -\mathbf{r}^{(t+1)} + \beta^{(t+1)} \mathbf{p}^{(t)}$$

$t \leftarrow t + 1$
end while

- Theorem: With CG algorithm

1. $\mathbf{r}^{(t)}$ are mutually orthogonal.
2. $\{\mathbf{r}^{(0)}, \dots, \mathbf{r}^{(t)}\}$ is contained in the *Krylov* subspace of degree t for $\mathbf{r}^{(0)}$, denoted by

$$\mathcal{K}(\mathbf{r}^{(0)}; t) = \text{span}\{\mathbf{r}^{(0)}, \mathbf{A}\mathbf{r}^{(0)}, \mathbf{A}^2\mathbf{r}^{(0)}, \dots, \mathbf{A}^t\mathbf{r}^{(0)}\}.$$

3. $\{\mathbf{p}^{(0)}, \dots, \mathbf{p}^{(t)}\}$ is contained in $\mathcal{K}(\mathbf{r}^{(0)}; t)$.
4. $\mathbf{p}^{(0)}, \dots, \mathbf{p}^{(t)}$ are conjugate wrt \mathbf{A} .

The iterates $\mathbf{x}^{(t)}$ converge to the solution in at most n steps.

- CG algorithm (economical version): saves one matrix-vector multiplication.

Given $\mathbf{x}^{(0)}$

Initialize: $\mathbf{r}^{(0)} \leftarrow \mathbf{A}\mathbf{x}^{(0)} - \mathbf{b}$, $\mathbf{p}^{(0)} \leftarrow -\mathbf{r}^{(0)}$, $t = 0$

while $\mathbf{r}^{(0)} \neq \mathbf{0}$ **do**

$$\alpha^{(t)} \leftarrow \frac{\mathbf{r}^{(t)\top} \mathbf{r}^{(t)}}{\mathbf{p}^{(t)\top} \mathbf{A} \mathbf{p}^{(t)}}$$

$$\mathbf{x}^{(t+1)} \leftarrow \mathbf{x}^{(t)} + \alpha^{(t)} \mathbf{p}^{(t)}$$

$$\mathbf{r}^{(t+1)} \leftarrow \mathbf{r}^{(t)} + \alpha^{(t)} \mathbf{A} \mathbf{p}^{(t)}$$

$$\beta^{(t+1)} \leftarrow \frac{\mathbf{r}^{(t+1)\top} \mathbf{r}^{(t+1)}}{\mathbf{r}^{(t)\top} \mathbf{r}^{(t)}}$$

$$\mathbf{p}^{(t+1)} \leftarrow -\mathbf{r}^{(t+1)} + \beta^{(t+1)} \mathbf{p}^{(t)}$$

$t \leftarrow t + 1$

end while

- Computation cost per iteration is one matrix vector multiplication: $\mathbf{A}\mathbf{p}^{(t)}$.

Consider PageRank problem, \mathbf{A} has dimension $n \approx 10^{10}$ but is highly structured (sparse + low rank). Each matrix vector multiplication takes $O(n)$.

- Theorem: If \mathbf{A} has r distinct eigenvalues, $\mathbf{x}^{(t)}$ converges to solution \mathbf{x}^* in at most r steps.

25 Lecture 25, Nov 25

Announcements

- HW8 (ranking MLB teams) due today. Submit hardcopy and email code (LastFirstHW8.R or LastFirstHW8.Rmd).
- HW9 (simulation project) due Dec 9 @ 11A.
- No office hours this Thu and Fri (Thanksgiving).

Last time

- Quasi-Newton method.
 - DFP: keep rank-2 updating inverse of approximate Hessian
 - BFGS: keep rank-2 updating approximate Hessian inverse
- (Linear) conjugate gradient method for solving linear equation $\mathbf{Ax} = \mathbf{b}$.

Today

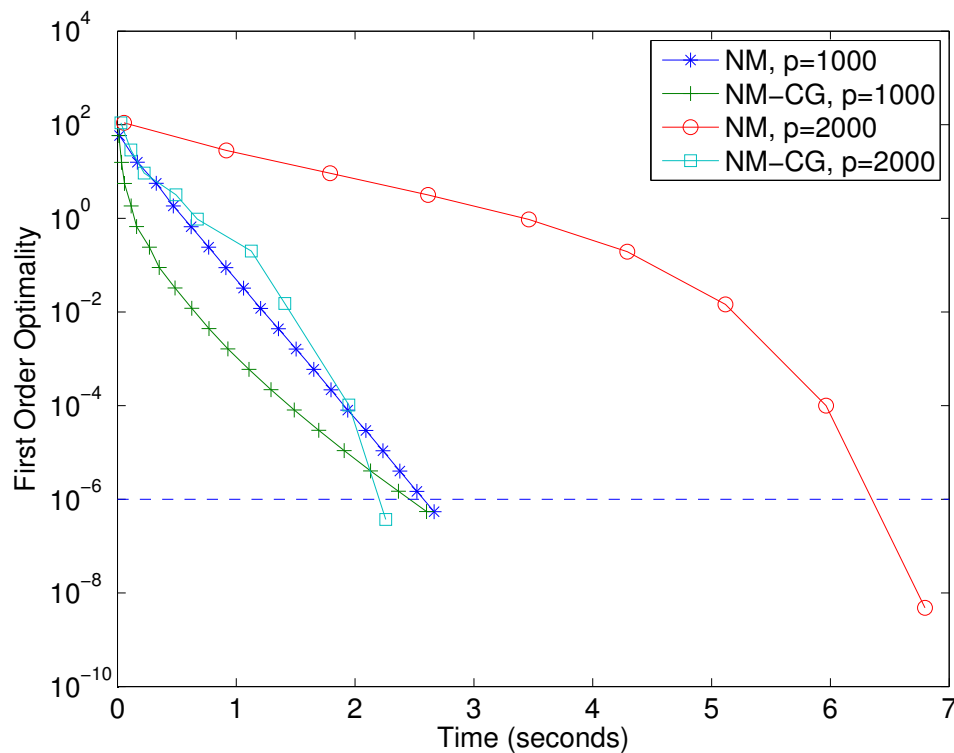
- Preconditioned conjugate gradient.
- (Nonlinear) conjugate gradient method.
- Concluding remarks on optimization.

Pre-conditioned conjugate gradient (PCG)

- Summary of conjugate gradient method for solving $\mathbf{Ax} = \mathbf{b}$ or equivalently minimizing $\frac{1}{2}\mathbf{x}^T \mathbf{Ax} - \mathbf{b}^T \mathbf{x}$:
 - Each iteration needs one matrix vector multiplication: $\mathbf{Ap}^{(t+1)}$. For structured \mathbf{A} , often $O(n)$ cost per iteration.
 - Guaranteed to converge in n steps.

- One example of using CG. Consider the Newton method implemented in HWS for ranking teams. Suppose the number of teams/players p is huge (e.g. online Chess). Most likely the Hessian is sparse. Then CG can be used to compute the Newton direction.

Simulation setup in the following figure: $p = 1000$ or 2000 , strengths of team are $\lambda_i = i + (p/10)$, each player competes with about $p/20$ opponents.



- Two important bounds for conjugate gradient algorithm:

Let $\lambda_1 \leq \dots \leq \lambda_n$ be the ordered eigenvalues of a pd \mathbf{A} .

$$\|\mathbf{x}^{(t+1)} - \mathbf{x}^*\|_{\mathbf{A}}^2 \leq \left(\frac{\lambda_{n-t} - \lambda_1}{\lambda_{n-t} + \lambda_1} \right)^2 \|\mathbf{x}^{(0)} - \mathbf{x}^*\|_{\mathbf{A}}^2$$

$$\|\mathbf{x}^{(t+1)} - \mathbf{x}^*\|_{\mathbf{A}}^2 \leq 2 \left(\frac{\sqrt{\kappa(\mathbf{A})} - 1}{\sqrt{\kappa(\mathbf{A})} + 1} \right)^t \|\mathbf{x}^{(0)} - \mathbf{x}^*\|_{\mathbf{A}}^2,$$

where $\kappa(\mathbf{A}) = \lambda_n/\lambda_1$ is the condition number of \mathbf{A} .




Figure 5.3 Two clusters of eigenvalues.




Figure 5.4 Performance of the conjugate gradient method on (a) a problem in which five of the eigenvalues are large and the remainder are clustered near 1, and (b) a matrix with uniformly distributed eigenvalues.

- Messages:
 - Roughly speaking, if the eigenvalues of \mathbf{A} occur in r distinct clusters, the CG iterates will *approximately* solve the problem after $O(r)$ steps.
 - \mathbf{A} with a small condition number ($\lambda_1 \approx \lambda_n$) converges fast.
- *Pre-conditioning*: Change of variables $\hat{\mathbf{x}} = \mathbf{C}\mathbf{x}$ via a nonsingular \mathbf{C} and solve

$$(\mathbf{C}^{-T}\mathbf{A}\mathbf{C}^{-1})\hat{\mathbf{x}} = \mathbf{C}^{-T}\mathbf{b}.$$

Choose \mathbf{C} such that

- $\mathbf{C}^{-T}\mathbf{A}\mathbf{C}^{-1}$ has small condition number, or
- $\mathbf{C}^{-T}\mathbf{A}\mathbf{C}^{-1}$ has clustered eigenvalues
- Inexpensive solution of $\mathbf{C}^T\mathbf{C}\mathbf{y} = \mathbf{r}$

- Preconditioned CG does not make use of \mathbf{C} explicitly, but rather the matrix $\mathbf{M} = \mathbf{C}^T \mathbf{C}$.
- Preconditioned CG (PCG) algorithm:

Given $\mathbf{x}^{(0)}$, pre-conditioner \mathbf{M}
 $\mathbf{r}^{(0)} \leftarrow \mathbf{A}\mathbf{x}^{(0)} - \mathbf{b}$
 solve $\mathbf{M}\mathbf{y}^{(0)} = \mathbf{r}^{(0)}$ for $\mathbf{y}^{(0)}$
 $\mathbf{p}^{(0)} \leftarrow -\mathbf{r}^{(0)}$, $t = 0$
while $\mathbf{r}^{(0)} \neq \mathbf{0}$ **do**
 $\alpha^{(t)} \leftarrow \frac{\mathbf{r}^{(t)\top} \mathbf{y}^{(t)}}{\mathbf{p}^{(t)\top} \mathbf{A} \mathbf{p}^{(t)}}$
 $\mathbf{x}^{(t+1)} \leftarrow \mathbf{x}^{(t)} + \alpha^{(t)} \mathbf{p}^{(t)}$
 $\mathbf{r}^{(t+1)} \leftarrow \mathbf{r}^{(t)} + \alpha^{(t)} \mathbf{A} \mathbf{p}^{(t)}$
 Solve $\mathbf{M}\mathbf{y}^{(t+1)} \leftarrow \mathbf{r}^{(t+1)}$ for $\mathbf{y}^{(t+1)}$
 $\beta^{(t+1)} \leftarrow \frac{\mathbf{r}^{(t+1)\top} \mathbf{y}^{(t+1)}}{\mathbf{r}^{(t)\top} \mathbf{r}^{(t)}}$
 $\mathbf{p}^{(t+1)} \leftarrow -\mathbf{y}^{(t+1)} + \beta^{(t+1)} \mathbf{p}^{(t)}$
 $t \leftarrow t + 1$
end while

Remark: Only extra cost in the pre-conditioned CG algorithm is the need to solve the linear system $\mathbf{M}\mathbf{y} = \mathbf{r}$.

- Pre-conditioning is more like an art than science. Some choices include
 - Incomplete Cholesky. $\mathbf{A} \approx \tilde{\mathbf{L}}\tilde{\mathbf{L}}^T$, where $\tilde{\mathbf{L}}$ is a sparse approximate Cholesky factor. Then $\tilde{\mathbf{L}}^{-1}\mathbf{A}\tilde{\mathbf{L}}^{-T} \approx \mathbf{I}$ (perfectly conditioned) and $\mathbf{M}\mathbf{y} = \tilde{\mathbf{L}}\tilde{\mathbf{L}}^T\mathbf{y} = \mathbf{r}$ is easy to solve.
 - Banded pre-conditioners.
 - Choose \mathbf{M} as a coarsened version of \mathbf{A} .
 - *Subject knowledge*. Knowledge about the structure and origin of a problem is often the key to devising efficient pre-conditioner. For example, see recent work of Stein et al. (2012) for pre-conditioning large covariance matrices. <http://epubs.siam.org/doi/abs/10.1137/110834469>

More buzzwords and softwares

Here are a few variants of CG that we should at least know the names and what they are for (so we can Google later in need).

- MINRES (minimum residual method): symmetric indefinite \mathbf{A} .
- Bi-CG (bi-conjugate gradient): unsymmetric \mathbf{A} .
- Bi-CGSTAB (Bi-CG stabilized): improved version of Bi-CG.
- GMRES (generalized minimum residual method): current *de facto* method for unsymmetric \mathbf{A} . E.g., PageRank problem.
- Lanczos method: top eigen-pairs of a large symmetric matrix.
- Arnoldi method: top eigen-pairs of a large unsymmetric matrix.
- Lanczos bidiagonalization algorithm: top singular triples of large matrix.

Remark: For Lanczos/Arnoldi methods, the critical computation is still matrix vector multiplication $\mathbf{A}\mathbf{v}$.

Softwares:

- MATLAB:
 - Iterative methods for solving linear equations:
`pcg`, `bicg`, `bicgstab`, `gmres`, ...
 - Iterative methods for top eigen-pairs and singular pairs:
`eigs`, `svds`, ...
 - Pre-conditioner:
`cholinc`, `luinc`, ...

Get familiar with the reverse communication interface (RCI) for utilizing iterative solvers:

```
x = gmres(A, b)
x = gmres(@Afun, b)
eigs(A)
eigs(@Afun)
```

- Consider the PageRank problem. We want to find the top left eigenvector of the transition matrix

$$\mathbf{P} = p\mathbf{R}^+ \mathbf{A} + \mathbf{z}\mathbf{1}_n^\top,$$

where $\mathbf{R} = \text{diag}(r_1, \dots, r_n)$ and $z_j = (1 - p)/n$ if $r_i > 0$ and $1/n$ if $r_i = 0$. Size of \mathbf{P} can be huge: $n \approx 40$ billion web pages. How to call the `gmres` or `eigs` function?

- R: Try Google and good luck ...

(Nonlinear) Conjugate gradient method

- *Linear* conjugate gradient method is for solving linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$, or equivalently, minimizing $\frac{1}{2}\mathbf{x}^T \mathbf{A}\mathbf{x} - \mathbf{b}^T \mathbf{x}$.
- *Nonlinear* conjugate gradient is for nonlinear optimization

$$\text{minimize } f(\mathbf{x}).$$

- History: Fletcher and Reeves in 60s.
- Fletcher-Reeves CG for nonlinear minimization:

```

Given  $\mathbf{x}^{(0)}$ 
Evaluate  $\nabla f^{(0)} = \nabla f(\mathbf{x}^{(0)})$ 
Set  $\mathbf{p}^{(0)} \leftarrow -\nabla f^{(0)}$ ,  $t \leftarrow 0$ 
while  $\nabla f^{(t)} \neq \mathbf{0}$  do
  Compute  $\alpha^{(t)}$  and set  $\mathbf{x}^{(t+1)} \leftarrow \mathbf{x}^{(t)} + \alpha^{(t)}\mathbf{p}^{(t)}$ 
  Evaluate  $\nabla f^{(t+1)} = \nabla f(\mathbf{x}^{(t+1)})$ 
   $\beta^{(t+1)} \leftarrow \frac{df^{(t+1)}\nabla f^{(t+1)}}{df^{(t)}\nabla f^{(t)}}$ 
   $\mathbf{p}^{(t+1)} \leftarrow -\nabla f^{(t+1)} + \beta^{(t+1)}\mathbf{p}^{(t)}$ 
   $t \leftarrow t + 1$ 
end while

```

- Most cost is evaluation of objective function and its gradient. No matrix operations are needed. Appealing for large nonlinear optimization problems.
- Line search (choose $\alpha^{(t)}$) is necessary to get a descending algorithm.

Concluding remarks on optimization

- Nonlinear optimization algorithms

Algorithm	Convergence Rate	Per-iteration Cost	Example
Newton	quadratic	high, usually $O(np^2) + O(p^3)$	GLM with canonical link
Fisher Scoring	super-linear	high, usually $O(np^2) + O(p^3)$	GLM with non-canonical link
Gauss-Newton	super-linear	high, usually $O(np^2) + O(p^3)$	nonlinear GLM
Quasi-Newton	super-linear	moderate, usually $O(np) + O(p^2)$	
Conjugate gradient	super-linear	moderate, usually $O(np) + O(p^2)$	
Coordinate descent	linear	low	
Steepest descent	linear	low	
EM/MM	linear	low	


There is simply no such thing as a universal ‘gold standard’ when it comes to algorithms.

Unknown Reviewer

- *Problem specific analysis* is critical for developing a successful optimization algorithm.

E.g., I don’t think any black-box procedure can beat our safe-guarded Newton’s method for the Dirichlet-Multinomial MLE problem (HW6/7) in terms of efficiency.

- Use “black-box” for
 - numerical linear algebra
 - least squares
 - convex programming (TODO in ST790-003 *Advanced Statistical Computing*). Current “technology” (Cplex, Gurobi, Mosek, cvx, Matlab, ..., R is not here ☺) can deal with problem with up to $10^3 \sim 10^4$ variables and constraints or even more with structure.



For example, the optimization problem in HW8 (ranking MLB teams)

$$\text{minimize} \quad \prod_{i,j} \left(\frac{\gamma_i}{\gamma_i + \gamma_j} \right)^{-y_{ij}} = \prod_{i,j} \gamma_i^{-y_{ij}} (\gamma_i + \gamma_j)^{y_{ij}}$$

is recognized as a standard geometric programming (GP) problem. Using the open source convex optimization software CVX (Grant and Boyd, 2012) amounts to only 6 lines of MATLAB code

```
p = size(Y, 1);
[rowidx, colidx, yvec] = find(Y);
cvx_begin gp
    variable gamma(p)
    minimize prod(((gamma(rowidx) + gamma(colidx)) ./ gamma(rowidx)) .^ yvec)
cvx_end
```

- First order methods (EM/MM, CD, steepest descent) is easier for parallel computing.

Algorithm development goes hand in glove with hardware advancement.

Hua

- Reference books:

- *Numerical Optimization* (Nocedal and Wright, 2006)
- *Convex Optimization* (Boyd and Vandenberghe, 2004)
- *The EM algorithm and Extensions* (McLachlan and Krishnan, 2008)
- *Numerical Analysis for Statisticians* (Lange, 2010)

26 Lecture 26, Dec 2

Announcements

- HW8 (ranking MLB teams) returned. Feedback:
 - Bradley-Terry model. KL 12.6 or David Hunter’s paper.
 - Non-concavity of the log-likelihood in λ parameterization. Enough to find a simple counter-example. However the negative log-likelihood is an example of geometric program, a branch of convex programming.
 - Vectorization of MM update: `outer` function.
 - Concavity of the log-likelihood function in γ parameterization. “log-sum-exp” terms are convex.
 - Implementation of Newton’s method. Hessian is singular due to identifiability. Setting γ_1 .
 - Problem structure: sparsity in large league.
 - Check David Hunter’s code for taking advantage of sparsity.

Sketch of solution: <http://hua-zhou.github.io/teaching/st758-2014fall/hw08sol.html>

- HW9 (simulation project) due Dec 9 @ 11A.
- FAQs on HW9 (simulation project)
- Regular office hours this week: Tue (Hua), Thu (Hua) and Fri (William). No office hours next week.
- **Course evaluation!**: <https://classeval.ncsu.edu/>

Last time

- Pre-conditioning for conjugate gradient (PCG) method.
- Nonlinear conjugate gradient for optimization.
- Concluding remarks on optimization.

Today

- Introduction to Markov chains and MCMC.
- Fast algorithms: sorting, FFT.
- Take home messages.

Introduction to MCMC

Some topics I'll briefly talk about.

- History of Markov chain
- History of Monte Carlo and birth of MCMC
- Convergence rate of Markov chain

Markov chains

- Markov chain is a stochastic process X_0, X_1, X_2, \dots with the *Markov property*

$$\mathbf{P}(X_{t+1}|X_t, X_{t-1}, \dots, X_0) = \mathbf{P}(X_{t+1}|X_t).$$

Given current state, the future is independent of the past.

- Stochastic analog of ordinary differential equations.

$$\frac{dx(t)}{dt} = F(x(t))$$
$$\mathcal{L}(X_{t+1}|X_t = x) = K(x, \cdot)$$

- Notations (for discrete time, finite Markov chains)
 - a finite state space \mathcal{X}
 - transition matrix $K(x, y)$, $x, y \in \mathcal{X}$
 - stationary distribution π on \mathcal{X} , defined as a probability vector that satisfies

$$\pi^T K = \pi^T.$$

Existence of such π is guaranteed by the Perron-Frobenius theorem.

- $K^l(x, y)$ denotes the l -step transition probabilities
- Markov chains in early years:
 - Fermat and Pascal (circa 1654): Gambler’s ruin.
 - Bernoulli (1769) and Laplace (1812): Urn model.
 - I. J. Bienaymé (1845), and later Sir Francis Galton and Watson (*Educational Times*, 1873): Branching process.
 - Paul and Tatiana Ehrenfest (1906): Statistical physics.
 - Poincaré (1912): Card shuffling. *Calcul des Probabilités*



- Markov’s contribution
 - * Markov, A. A. (1906)
Extension of the law of large numbers to dependent quantities [in Russian], *Izv. Fiz.-Matem. Obsch. Kazan Univ. (2nd Ser.)*
 - * St. Petersburg School vs Moscow School.
 - * Example: 20,000 letters in Pushkin’s *Eugene Onegin*.

$$\begin{array}{r}
 \text{vowel} \\
 \text{consonant}
 \end{array}
 \begin{pmatrix}
 \text{vowel} & \text{consonant} \\
 0.128 & 0.872 \\
 0.663 & 0.337
 \end{pmatrix}
 \begin{pmatrix}
 \pi \\
 0.432 \\
 0.568
 \end{pmatrix}$$

- * See Seneta’s interesting account on the history (Seneta, 1996)
- Some other leading pioneers: Kolmogorov, Fréchet, and Doeblin.
- Why are Markov chains important in statistics?

- Modeling tool.
E.g., Marc Coram’s “jail message” example: Markov model for letter sequence; PageRank’s (imaginary) random web surfer; ...
- A methodology that revolutionized statistics: Markov Chain Monte Carlo (MCMC).
 - * Metropolis, Rosenbluth, Rosenbluth, Teller and Teller (1953)
Equation of state calculation by fast computing machines, *The Journal of Chemical Physics*, 21, 1087–1092.
 - * Want to sample from a distribution $\pi(x) \propto f(x)$.
Metropolis algorithm constructs a Markov chain that converges to π .
 - * Marc Coram’s “jail message” example.

Markov chain Monte Carlo (MCMC)

- *Monte Carlo method* is a generic name for “computational algorithms that rely on repeated random sampling to obtain numerical results”. They are in contrast to the deterministic algorithms. They have wide applications in
 - integration
 - drawing sample from a distribution. E.g., HW9 (simulation study).
 - optimization (simulated annealing). E.g., Marc Coram’s “jail message” example, traveling salesman, Soduku, ...
- Monte Carlo in early years



- Stanislaw Ulam conceived it in 1946 while playing solitaire in hospital bed. He wanted to know the probability of getting a perfect solitaire hand, and wondered whether computers can help answer this.
- John von Neumann was intrigued by the idea and developed a way to generate pseudorandom numbers (inversion, importance sampling, acceptance-rejection sampling) on electronic digital computers (ENIAC) to realize Ulam's idea, for the neutron fission and diffusion problem.

ANOTHER VON NEUMANN LETTER

Fig. 3. In this 1947 letter to Stan Ulam, von Neumann discusses two methods for generating the nonuniform distributions of random numbers needed in the Monte Carlo method. The second paragraph summarizes the inverse-function approach in which (x') represents the uniform distribution and (ξ') the desired nonuniform distribution. The rest of the letter describes an alternative approach based on two uniform and independent distributions: (x') and (y') . In this latter approach a value x' from the first set is accepted when a value y' from the second set satisfies the condition $y' \leq f(x')$, where $f(\xi')$ is the density of the desired distribution function. (It should be noted that in von Neumann's example for forming the random pairs $(\xi = \sin x$ and $\eta = \cos x$, he probably meant to say that x is equidistributed between 0 and 360 degrees (rather than "300"). Also, his notation for the tangent function is tg , so that the second set of equations for x and y are just half-angle ($y = x/2$) trigonometric identities.)

digits 0, . . . , 9	Replace ξ, η by	ξ, η
0	+	$-\xi, \eta$
1	+	ξ, η
2	+	$-\xi, -\eta$
3	+	$\xi, -\eta$
4	+	$-\xi, \eta$
5	+	ξ, η
6	+	$-\xi, -\eta$
7	+	ξ, η
8	+	$-\xi, -\eta$
9	+	ξ, η

Reject this digit
 Now $t = \xi \eta$, $t^2 = \xi^2 \eta^2 = 4/5$, lies between 0 and 1, and its distribution function is $\frac{t^2}{1-t^2}$. Hence one may pick pairs of numbers ξ, η both (independently) equidistributed between 0 and 1, and then use t for $(1+t^2) \leq 1$
 reject t for $(1+t^2) > 1$
 form $\frac{t}{1+t^2}$ and $\frac{\eta}{1+t^2}$ then step
 Of course, the first pair . . . unless a divider, but the method may still be worth keeping in mind, especially when the entire is available.

with best regards from house to house.
 Yours, as ever,
 John von Neumann

JvN:JH

THE INSTITUTE FOR ADVANCED STUDY
 SCHOOL OF MATHEMATICS
 PRINCETON, NEW JERSEY

May 21, 1947

Mr. Stan Ulam
 Post Office Box 1663
 Radio Fe
 New Mexico

Dear Stan

Thanks for your letter of the 18th. I need not tell you that Elard and I are looking forward to the trip and visit at Los Alamos this summer. I have already received the necessary papers from Carson Mark. I filled out and returned mine yesterday; Elard's will follow today.

I am very glad that preparations for the random numbers work are to begin soon. In this connection, I would like to mention this: Assume $\xi, \eta = (\xi^1, \eta^1), (\xi^2, \eta^2), \dots$. Assume that you want one with the distribution function (density) $f(\xi)$ and $g(\eta)$. One way to form it is to form the cumulative distribution functions $F(\xi) = \int_{-\infty}^{\xi} f(x) dx$ with this $F(\xi) = \int_{-\infty}^{\xi} f(x) dx$, and to form F^{-1} by some approximate polynomial. This is, as I see, the method that you have in mind.

An alternative, which works if f and all values of $f(\xi)$ lie in $[0, 1]$, is this: form pairs ξ, η and use or reject ξ, η according to whether $y' \leq f(x')$ or not. In the first case, put $\xi^d = x'$ in the second case form no ξ^d at that step.

The second method may occasionally be better than the first one. In some cases combinations of both may be best; e.g., form random pairs $\xi = \sin x, \eta = \cos x$ with x equidistributed between 0° and 360° . The obvious way consists of using the $\sin - \cos$ values (with interpolation). This is clearly closely related to the first method. This is an alternative procedure:
 Put $\xi = \frac{2-t}{1+t^2}, \eta = \frac{1-t^2}{1+t^2}, t = \xi \eta$

with y (which is $\frac{2-t}{1+t^2}$) equidistributed between 0° and 180° . Restrict y to 0° to 45° . Then the ξ, η will have to be replaced randomly by η, ξ and again by $\pm \xi, \pm \eta$. This can be done by using random digits 0, . . . , 9. It is also feasible with

185

- Ulam and von Neumann, working on Manhattan Project, used the code name *Monte Carlo*. It is a casino in Monaco where Ulam's uncle frequented.
- Nicholas Metropolis, fascinated by the Monte Carlo idea too, designed and built computing devices (MANIAC) to handle such calculations. His paper with Ulam in JASA (Metropolis and Ulam, 1949) formed the basis of modern sequential Monte Carlo methods.
- Birth of MCMC (Metropolis et al., 1953):
 - * Metropolis, Rosenbluth, Rosenbluth, Teller and Teller (1953) Equation of state calculation by fast computing machines, *The Journal of Chemical Physics*, 21, 1087-1092.
 - * Want to sample from a distribution $\pi(x) \propto f(x)$.
 Metropolis algorithm constructs a Markov chain that converges to π .
 - * *Metropolis chain*: From current state x , generate a new state x' (from a


proposal distribution $p(x, x')$ such that $p(x, x') = p(x', x)$ and accept x' with probability $\min\left\{\frac{f(x')}{f(x)}, 1\right\}$.

Fact: Metropolis chain has π as stationary distribution.

* Marc Coram's "jail message" example.



- Given π , generic ways to construct a Markov chain K that has π as stationary distribution:
 - Metropolis algorithm: (Metropolis et al., 1953)
 - Hastings algorithm: (Hastings, 1970)
 - Gibbs sampler: Glauber dynamics (1963), Tanner and Wong (1987), Gelfand and Smith (1990)
- See KL Ch25-27 and JM Ch13 for a general introduction, or take a Bayesian course. A comprehensive textbook is (Robert and Casella, 2004).





Convergence rate of Markov chains

- Classical result: For finite, irreducible, and aperiodic Markov chains

$$\lim_{l \rightarrow \infty} K^l(x, y) = \pi(y).$$

In practice, we often want to know *how many* steps to make the difference between $K^l(x, \cdot)$ and π small?

- Example: “How many shuffles do I need to do to mix a deck of 52 cards?”



Consider *riffle shuffle*. Gilbert-Shannon-Reeds model: binomial(52, 0.5) cut + cards drop according to probability $L/(L + R)$, where L and R are the number of cards in the left and right hand respectively.

- Example: "How long do I need to run my Gibbs sampler?"

Consider the Beta-Binomial Gibbs sampler

- Likelihood $f(x | p) \sim \text{Bin}(n, p)$ and prior $\pi(p) \sim \text{Beta}(\alpha, \beta)$
 - Want to sample from joint density $f(x, p) = f(x | p)\pi(p)$
 - *Gibbs sampler*: Repeat the following
 - * Sample x from $\text{Bin}(n, p)$
 - * Sample p from $\text{Beta}(x + \alpha, n - x + \beta)$
 - $(X_l, p_l)_{l \geq 1}$ form a Markov chain on $\{0, 1, \dots, n\} \times [0, 1]$
 - Let $\tilde{K}(x, p; x', p')$ be the transition density
 - How many steps (obviously depending on n, α, β) does this Markov chain converge to the stationary distribution?
- In a typical Bayesian course, we learn many convergence diagnostics that are often heuristic. Is there anything rigorous we can say about convergence rate of Markov chains?
 - Distances between distributions:

– *Total variation distance*:

$$\begin{aligned} \|\mu - \pi\|_{\text{TV}} &= \frac{1}{2} \sum_{x \in \mathcal{X}} |\mu(x) - \pi(x)| \\ &= \max_{A \subset \mathcal{X}} |\mu(A) - \pi(A)| \\ &= \frac{1}{2} \sup_{\|f\|_{\infty} \leq 1} |\mu(f) - \pi(f)|. \end{aligned}$$

– L^p distance wrt π :

Let $f(x) = \frac{\mu(x)}{\pi(x)}$ and $g(x) = \frac{\nu(x)}{\pi(x)}$. For $1 \leq p < \infty$,

$$d_{\pi,p}(\mu, \nu) = \|f - g\|_{L^p(\pi)} = \left(\sum_{x \in \mathcal{X}} |f(x) - g(x)|^p \pi(x) \right)^{1/p}$$

- The usual limit theorems are useless in practice:

"There exist constants $A, B > 0$, $\rho \in (0, 1)$ such that $\|K^l(x, \cdot) - \pi\|_{\text{TV}} \leq A\rho^{Bl}$."

A , B , and ρ are some mysterious constants.

Can we get some more quantitative, useful bounds?

- Cutoff phenomenon (Diaconis, 1996)




FIG. 1. The cutoff phenomenon for repeated riffle shuffles of $n = 52$ cards.

- Riffle shuffle.

l	1	2	3	4	5	6	7	8	9	10
$\ K^l - \pi\ _{TV}$	1.000	1.000	1.000	1.000	0.924	0.624	0.312	0.161	0.083	0.041

- Many more examples: Ehrenfest chain, random transposition, Gibbs sampler, ...

“Cutoff phenomenon for XXX chain.”

- Not every Markov chain has a cutoff.

A chain without cutoff: simple random walk on the integers mod n .

- Generic methods for studying convergence rates:
 - Algebraic methods (spectral analysis) - E.g, random walks on groups (shuffling cards), some Gibbs samplers, ...
 - Analytic methods - Geometric Inequalities.
 - Probabilistic methods - Coupling and strong stationary times.
- Algebraic method

– *Reversible* Markov chains.

If π is a probability distribution on \mathcal{X} and

$$\pi(x)K(x, y) = \pi(y)K(y, x), \quad \text{for all } x, y \in \mathcal{X},$$

then π is the unique stationary distribution of K . E.g., Metropolis chain.

– K operates on $L^2(\pi) = \{f : \mathcal{X} \mapsto \mathbb{R}, \mathbf{E}_\pi[f^2] < \infty\}$ by

$$Kf(x) = \sum_{y \in \mathcal{X}} K(x, y)f(y).$$

– Reversibility of K is equivalent to *self-adjointness* of the operator K .

– By standard spectral theorem for self-adjoint operators, K has eigenvalues $1 = \beta_0 \geq \beta_1 \geq \dots \geq \beta_{|\mathcal{X}|-1} \geq -1$ with (right) eigenfunctions $\{\phi_0, \dots, \phi_{|\mathcal{X}|-1}\}$ that are orthonormal on $L^2(\pi)$

$$\langle \phi_i, \phi_j \rangle_{L^2(\pi)} = \sum_{x \in \mathcal{X}} \phi_i(x)\phi_j(x)\pi(x) = 1_{\{i=j\}}.$$

– If we know all the spectral information (lucky!), then

$$d_{\pi,2}^2(K^l(x, \cdot), \pi) = \sum_{i=1}^{|\mathcal{X}|-1} \beta_i^{2l} \phi_i^2(x).$$

– Usually the upper bound is tight.

$$\|K^l(x, \cdot) - \pi\|_{TV} \leq \frac{1}{2}d_{\pi,2}(K^l(x, \cdot), \pi)$$

– When are we lucky? In presence of *symmetry*.

E.g., for random walks on groups, we only need eigenvalues, which can be derived from the irreducible representations of the group.

Definite reference for this topic is the book Diaconis (1988)

- Algebraic method for analyzing the baby Gibbs sampler

– The joint chain $\tilde{K}(x, p; x', p')$ is *irreversible*.

– The x -marginal chain $K(x, x')$ is *reversible*, with $m \sim \text{Beta-Bin}(n, \alpha, \beta)$ as stationary distribution. And

$$\|K_x^l - m\|_{TV} \leq \|\tilde{K}_{x,p}^l - f\|_{TV} \leq \|K_x^{l-1} - m\|_{TV}.$$

Thus sufficient to study convergence rate of the x -marginal chain.

- Some analysis (Diaconis et al., 2008) shows that K has
 - * eigenvalues: $\beta_0 = 1, \beta_j = n_{[j]}/(n + \alpha + \beta)_{(j)}, j = 1, \dots, n.$
 - * eigen-functions: ϕ_j are the *Hahn polynomials*.
- Doing the summation gives

$$0.5\beta_1^l \leq \|\tilde{K}_{0,p}^l - f\|_{TV} \leq 3\beta_1^{l-1/2}.$$

- Cutoff phenomenon: $\frac{n+\alpha+\beta}{2(\alpha+\beta)}$ steps are necessary and sufficient for convergence.
- Similar analysis can be carried out for all following conjugate pairs (see KL 27.9)

TABLE 26.1. Conjugate Pairs

Likelihood	Density	Prior	Density
Binomial	$\binom{n}{x} p^x (1-p)^{n-x}$	Beta	$\frac{1}{B(\alpha, \beta)} p^{\alpha-1} (1-p)^{\beta-1}$
Poisson	$\frac{\lambda^x}{x!} e^{-\lambda}$	Gamma	$\frac{\beta^\alpha \lambda^{\alpha-1}}{\Gamma(\alpha)} e^{-\beta\lambda}$
Geometric	$(1-p)^x p$	Beta	$\frac{1}{B(\alpha, \beta)} p^{\alpha-1} (1-p)^{\beta-1}$
Multinomial	$\binom{n}{x_1 \dots x_k} \prod_{i=1}^k p_i^{x_i}$	Dirichlet	$\frac{\Gamma(\sum_{i=1}^k \alpha_i)}{\prod_{i=1}^k \Gamma(\alpha_i)} \prod_{i=1}^k p_i^{\alpha_i-1}$
Normal	$\sqrt{\frac{\tau}{2\pi}} e^{-\tau(x-\mu)^2/2}$	Normal	$\sqrt{\frac{\omega}{2\pi}} e^{-\omega(\mu-\theta)^2/2}$
Normal	$\sqrt{\frac{\tau}{2\pi}} e^{-\tau(x-\mu)^2/2}$	Gamma	$\frac{\beta^\alpha \tau^{\alpha-1}}{\Gamma(\alpha)} e^{-\beta\tau}$
Exponential	$\lambda e^{-\lambda x}$	Gamma	$\frac{\beta^\alpha \lambda^{\alpha-1}}{\Gamma(\alpha)} e^{-\beta\lambda}$

- Analytic method

- Upper bound through spectral gap

$$4\|K^l(x, \cdot) - \pi\|_{TV} \leq d_{\pi,2}^2(K^l(x, \cdot), \pi) \leq \frac{1}{\pi(x)} \beta_*^{2l},$$

where $\beta_* = \max\{|\beta_1|, |\beta_{|\mathcal{X}|-1}|\}$.

- Use some geometric inequalities to bound β_* .
- Where to look up the material? Lecture notes by Saloff-Coste (1997).
- (K, π) reversible on \mathcal{X} (finite). Dirichlet form

$$\epsilon(f, g) = \langle (I - K)f, g \rangle.$$

Fact

$$\epsilon(f, f) = \frac{1}{2} \sum_{x,y} [f(x) - f(y)]^2 \pi(x) K(x, y).$$

– Lemma

$$1 - \beta_1 = \min_{f \neq 1} \frac{\epsilon(f, f)}{\text{var}(f)}$$
$$1 - \beta_{|\mathcal{X}|-1} = \max_{f \neq 0} \frac{\epsilon(f, f)}{\text{var}(f)}.$$

– Definition: Poincaré inequality

$$\text{var}(f) \leq A\epsilon(f, f), \quad \text{for all } f \in L^2(\pi).$$

– We can bound β_1 by finding A

$$\beta_1 \leq 1 - \frac{1}{A}.$$

– Theorem (Poincaré Ineq for Markov Chains (Diaconis and Stroock, 1991)):

$$\beta_1 \leq 1 - \frac{1}{A}, \quad \text{where } A = \max_e \frac{1}{Q(e)} \sum_{\gamma_{x,y} \ni e} |\gamma_{x,y}| \pi(x) \pi(y).$$

– Apply Poincaré inequality to the baby Gibbs sampler gives a horrible bound (something *exponential* ...)

– Other geometric inequalities: Cheeger, Nash, Log-Soblov inequalities (in a series of papers by Diaconis and Saloff-Coste).

- Probabilistic methods

- By cleverness, we can get good bounds for convergence rate without using all those analytic methods.

- Good starting point is the book (Diaconis, 1988) and the unpublished book by Aldous and Fill (available on Aldous' website).

- *Coupling* - Wolfgang Doeblin.




Fig. 20: Cover of the DVD: "Wolfgang Doeblin. A mathematician rediscovered"



Fig. 21: The archive of the "Académie des Sciences"

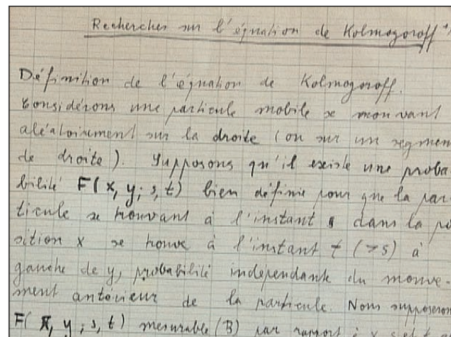


Fig.22: First page of the "pli cacheté no. 11668"

- For the baby Gibbs sampler, coupling gives an upper bound of order $n \ln n$ (off by $\ln n$).
- Yields useful bounds for hierarchical random effects model (Hobert and Geyer, 1998).
- Let's work on a simpler example: Borel's Shuffle (random to top, random to bottom, ...).
- *Coupling*: Two processes evolve until they are equal. *Coupling time* T . Coupling inequality

$$\|K_x^l - \pi\|_{TV} \leq \mathbf{P}(T > l).$$

- For Borel's shuffle, bound on coupon collector problem gives

$$\|K_x^l - \pi\|_{TV} \leq n \left(1 - \frac{1}{n}\right)^l.$$

Thus $l = n \ln n$ steps suffice.

- Summary
 - Keep Markov chains in your toolbox - useful for modeling, simulation, and combinatorial optimization.
 - Convergence rate of Markov chains is an interesting applied probability problem that often gives more insights into the chains.
 - A lot remains to be done for analyzing many MCMC algorithms being used.


Sorting algorithms (JM 14.3, KL 1.10)

- Applications: order statistics (median, quantiles), QQ-plot, multiple testing (sorting p-values), Wilcoxon rank-sum test, ...
- *Bubble sort*: Locate maximum and put on top; find maximum in the $(n - 1)$ list and put on the top second position; ...
 $O(n^2)$ average cost.
- Think about sorting *massive* data $n = 10^{12}$. On a teraflop computer. n^2 flops take 10^{12} seconds ≈ 31710 years, while $n \ln n$ flops take $10^{12} \log(10^{12})/10^{12} \approx 27$ seconds.
- Key idea: Divide and conquer.
- *Merge sort*: Recursively partition into two lists, sort them respectively, and then merge. $T(n) = 2T(n/2) + O(n)$. Solution is $T(n) = O(n \log_2 n)$.
- *Quick sort*: Randomly select a pivot element, split into 3 lists, and do some swaps so that the pivot is in the right position.

$$T(n) = \frac{1}{n} \sum_{j=1}^{n-1} [T(j-1) + T(n-j)] + n - 1 = \frac{2}{n} \sum_{j=1}^{n-1} T(j) + n - 1.$$

Solution is $O(n \ln n)$.

- Sorting is a well-trodden area in computer science. Mature functions/libraries in standard softwares. The “bible” on this topic is (of course) Knuth (2005).



Fast Fourier transform (FFT) (KL Chapter 20, JM 14.5)

- History: Cooley and Tukey (1965)

John Tukey: “bit”, box-plot, “learning from the experience of the others”, multiple comparison, FFT, ...

Tukey conceived the FFT algorithm during meetings with President JFK’s Science Advisory Committee. They need fast ways to analyze seismic waves to detect nuclear weapon tests in Soviet Union. Richard Garwin of IBM immediately realize the potential of this fast algorithm and referred Tukey to Cooley to implement it.

People also believe Gauss essentially used the same strategy when solving his least squares problem!

- Applications in statistics: convolution, time series, branching process, ...
- Consider two independent random variables on $\{0, 1, \dots, N - 1\}$:

$$X \sim \{p_0, \dots, p_{N-1}\}, \quad Y \sim \{q_0, \dots, q_{N-1}\}.$$

What’s the distribution of the sum $Z = X + Y$?

- $z_k = \sum_{j=0}^k p_j q_{k-j}$, $k = 0, \dots, 2N - 2$. $O(N^2)$ computation.
- Do DFT of both sequences, multiply together, and inverse DFT. $O(N \ln N)$ computation!

- Discrete Fourier transform (DFT) of a vector $\mathbf{x} \in \mathbb{R}^N$.

$$a_k = \sum_{j=0}^{N-1} w^{jk} x_j, \quad k = 0, \dots, N - 1.$$

where $w = e^{-2\pi\sqrt{-1}/N}$. Note w is an N -th root of 1. DFT is essentially matrix-vector multiplication

$$\mathbf{a}^\top = \mathbf{x}^\top \mathbf{W}, \quad \mathbf{W} = (w^{jk}),$$

which usually costs $O(N^2)$ flops.

- Suppose $N = N_1 N_2$. Index rewriting:
 - $j \leftarrow j_1 N_2 + j_2$ (fill out N_1 -by- N_2 matrix in row major),
 - $k \leftarrow k_2 N_1 + k_1$ (fill out N_1 -by- N_2 matrix in column major),
 - $j_1, k_1 \in \{0, \dots, N_1 - 1\}$, $j_2, k_2 \in \{0, \dots, N_2 - 1\}$.

Then

$$\begin{aligned} a_k &= a_{k_2 N_1 + k_1} = \sum_{j=0}^{N-1} w^{jk} x_j \\ &= \sum_{j_1=0}^{N_1-1} \sum_{j_2=0}^{N_2-1} w^{(j_1 N_2 + j_2)(k_2 N_1 + k_1)} x_{j_1 N_2 + j_2} \\ &= \sum_{j_1=0}^{N_1-1} \sum_{j_2=0}^{N_2-1} w^{j_1 k_1 N_2 + j_2 (k_2 N_1 + k_1)} x_{j_1 N_2 + j_2} \\ &= \sum_{j_2=0}^{N_2-1} w^{j_2 (k_2 N_1 + k_1)} \sum_{j_1=0}^{N_1-1} (w^{N_2})^{j_1 k_1} x_{j_1 N_2 + j_2} \\ &= \sum_{j_2=0}^{N_2-1} (w^{N_1})^{j_2 k_2} w^{j_2 k_1} \sum_{j_1=0}^{N_1-1} (w^{N_2})^{j_1 k_1} x_{j_1 N_2 + j_2}. \end{aligned}$$

- Essentially we need to do N_2 DFT of length- N_1 sequences and then do N_1 DFT of length- N_2 sequences. Total cost

$$T(N) = T(N_1 N_2) = N_2 T(N_1) + N_1 T(N_2).$$

Suppose N is a power of 2. Then $T(N) = (N/2)T(2) + 2T(N/2)$ and the solution is $T(N) = O(N \ln N)!$.

- Inverse DFT. \mathbf{W}^{-1} has entries w^{-jk}/N . Then

$$\mathbf{x}^\top = \mathbf{a}^\top \mathbf{W}^{-1}.$$

- Variants for prime N : still $O(N \ln N)$. But always a good idea to pad with zero to get N as a power of 2.
- Generalizations to 2D, 3D FFT available.
- Mature libraries/functions for both CPU and GPU.
- Galton-Watson process. Survival of families. Lotka data (Lotka, 1931a,b). Using 1920 census data, the *progeny generating function* for a white male

$$P(s) = .4982 + .2103s + .1270s^2 + .0730s^3 + .0418s^4 + .0241s^5 \\ + .0132s^6 + .0069s^7 + 0.0035s^8 + .0015s^9 + .0005s^{10}.$$

PGF for the first generation $P_1(s) = P(s)$. PGF for the second generation $P_2(s) = \sum_k p_k P(s)^k = P(P(s))$. In general, PGF for the i -th generation is $P_i(s) = P(\dots P(s))$ (i recursions).

Extinction probability of a family: $\lim_{i \rightarrow \infty} P_i(0) = P(\dots P(0)) = 0.88$, or solving for $P(s) = s$.

What if we want to know the distribution of the i -th generation? Extend the generating function P_i to unit circle $P_i(w^k) = \sum_j p_j w^{jk}$, $w = e^{-2\pi\sqrt{-1}/N}$, where $k = 0, \dots, N-1$ for N large. So $P_i(w^k)$ is the DFT of distribution p_j of i -th generation. Then apply inverse DFT to retrieve p_j . $O(N \log N)$ cost!

- Continuous-time branching process. Solve differential equation for $P_t(w^j)$ at any time t . Then apply inverse DFT.
- See JM 14.7 for more applications of FFT in statistics.

Take-home messages from this course

- Statistics, the science of *data analysis*, is the applied mathematics in the 21st century
 - Read the first and last few pages of Tukey (1962)'s *Future of data analysis*. <http://www.stat.ncsu.edu/people/zhou/courses/st810/notes/Tukey61FutureDataAnalysis.pdf>.
- *Big data* era: Challenges also mean opportunities for statisticians
 - methodology: big p
 - efficiency: big n and/or big p
 - memory: big n , distributed computing via MapReduce (Hadoop), online algorithms



- Being good at computing (*both* programming and algorithms) is a must for today's working statisticians.

Computers are incredibly fast, accurate, and stupid. Human beings are incredibly slow, inaccurate, and brilliant. Together they are powerful beyond imagination.

Albert Einstein

US (German-born) physicist (1879 - 1955)



- HPC (high performance computing) \neq abusing computers.
Always optimize your algorithms *as much as possible* before resorting to cluster computing resources.
- Coding
 - Prototyping: R, Matlab, Julia
 - A “real” programming language: C/C++, Fortran, Python
 - Scripting language: Python, Linux/Unix script, Perl, JavaScript
- Numerical linear algebra – building blocks of most computing we do. Use standard *libraries* (BLAS, LAPACK, ...)! Sparse linear algebra and iterative solvers such as conjugate gradient methods are critical for exploiting structure in big data.
- Optimization
 - *Convex programming* (LS, LP, QP, GP, SOCP, SDP). To do in ST790-003. Convex programming is becoming a *technology*, just like least squares (LS).
 - Specialized optimization algorithms for modern statistical learning problems. To do in ST790-003.
 - Generic nonlinear optimization tools: Newton, quasi-Newton, (nonlinear) conjugate gradient, ...
 - Specialized tools in statistics: EM/MM, Fisher scoring, Gauss-Newton, simulated annealing, ...

- Combinatorial optimization techniques: divide-and-conquer, dynamic programming, greedy algorithm, ...
- MCMC: take a Bayesian course!

References

- Baum, L. E., Petrie, T., Soules, G., and Weiss, N. (1970). A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *Ann. Math. Statist.*, 41:164–171.
- Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society, Ser. B*, 57:289–300.
- Björck, Å. (1996). *Numerical methods for least squares problems*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA.
- Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press, Cambridge.
- Cooley, J. W. and Tukey, J. W. (1965). An algorithm for the machine calculation of complex Fourier series. *Math. Comp.*, 19:297–301.
- Cox, D. R. (1972). Regression models and life-tables. *J. Roy. Statist. Soc. Ser. B*, 34:187–220. With discussion by F. Downton, Richard Peto, D. J. Bartholomew, D. V. Lindley, P. W. Glassborow, D. E. Barton, Susannah Howard, B. Benjamin, John J. Gart, L. D. Meshalkin, A. R. Kagan, M. Zelen, R. E. Barlow, Jack Kalbfleisch, R. L. Prentice and Norman Breslow, and a reply by D. R. Cox.
- Davidon, W. C. (1991). Variable metric method for minimization. *SIAM J. Optim.*, 1(1):1–17.
- Dempster, A., Laird, N., and Rubin, D. (1977a). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society Series B.*, 39(1-38).
- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977b). Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Statist. Soc. Ser. B*, 39(1):1–38. With discussion.
- Diaconis, P. (1988). *Group representations in probability and statistics*. Institute of Mathematical Statistics Lecture Notes—Monograph Series, 11. Institute of Mathematical Statistics, Hayward, CA.

- Diaconis, P. (1996). The cutoff phenomenon in finite Markov chains. *Proc. Nat. Acad. Sci. U.S.A.*, 93(4):1659–1664.
- Diaconis, P. (2009). The Markov chain Monte Carlo revolution. *Bull. Amer. Math. Soc. (N.S.)*, 46(2):179–205.
- Diaconis, P., Khare, K., and Saloff-Coste, L. (2008). Gibbs sampling, exponential families and orthogonal polynomials. *Statist. Sci.*, 23(2):151–200. With comments and a rejoinder by the authors.
- Diaconis, P. and Stroock, D. (1991). Geometric bounds for eigenvalues of Markov chains. *Ann. Appl. Probab.*, 1(1):36–61.
- Dickey, D. A. and Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *J. Amer. Statist. Assoc.*, 74(366, part 1):427–431.
- Efron, B. (1979). Bootstrap methods: another look at the jackknife. *Ann. Statist.*, 7(1):1–26.
- Efron, B. and Morris, C. (1973). Stein’s estimation rule and its competitors—an empirical Bayes approach. *J. Amer. Statist. Assoc.*, 68:117–130.
- Efron, B. and Morris, C. (1977). Stein’s paradox in statistics. *Scientific American*, 236(5):119–127.
- Gelfand, A. E. and Smith, A. F. M. (1990). Sampling-based approaches to calculating marginal densities. *J. Amer. Statist. Assoc.*, 85(410):398–409.
- Gentle, J. E. (2007). *Matrix Algebra*. Springer Texts in Statistics. Springer, New York. Theory, computations, and applications in statistics.
- Golub, G. H. and Van Loan, C. F. (1996). *Matrix Computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, third edition.
- Golub, G. H. and Van Loan, C. F. (2013). *Matrix Computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, fourth edition.

- Grant, M. and Boyd, S. (2012). CVX: Matlab software for disciplined convex programming, version 2.0 beta. <http://cvxr.com/cvx>.
- Harville, D. A. (1997). *Matrix Algebra From a Statistician's Perspectives*. Springer-Verlag, New York.
- Hastings, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109.
- Hobert, J. P. and Geyer, C. J. (1998). Geometric ergodicity of Gibbs and block Gibbs samplers for a hierarchical random effects model. *J. Multivariate Anal.*, 67(2):414–430.
- Horn, R. A. and Johnson, C. R. (1985). *Matrix Analysis*. Cambridge University Press, Cambridge.
- Huber, P. J. (1994). Huge data sets. In *COMPSTAT 1994 (Vienna)*, pages 3–13. Physica, Heidelberg.
- Huber, P. J. (1996). Massive data sets workshop: The morning after. In *Massive Data Sets: Proceedings of a Workshop*, pages 169–184. National Academy Press, Washington.
- Kaplan, E. L. and Meier, P. (1958). Nonparametric estimation from incomplete observations. *J. Amer. Statist. Assoc.*, 53:457–481.
- Knuth, D. E. (2005). *The Art of Computer Programming. Vol. 1. Fasc. 1*. Addison-Wesley, Upper Saddle River, NJ. MMIX, a RISC computer for the new millennium.
- Lange, K. (2010). *Numerical Analysis for Statisticians*. Statistics and Computing. Springer, New York, second edition.
- Lange, K. and Carson, R. (1984). EM reconstruction algorithms for emission and transmission tomography. *J. Comput. Assist. Tomogr.*, 8(2):306–316.
- Lange, K., Hunter, D. R., and Yang, I. (2000). Optimization transfer using surrogate objective functions. *J. Comput. Graph. Statist.*, 9(1):1–59. With discussion, and a rejoinder by Hunter and Lange.

- Lawson, C. L. and Hanson, R. J. (1987). *Solving Least Squares Problems*. Classics in Applied Mathematics. Society for Industrial Mathematics, new edition edition.
- Lotka, A. (1931a). Population analysis - the extinction of families i. *J. Wash. Acad. Sci.*, 21:377–380.
- Lotka, A. (1931b). Population analysis - the extinction of families ii. *J. Wash. Acad. Sci.*, 21:453–459.
- Magnus, J. R. and Neudecker, H. (1999). *Matrix Differential Calculus with Applications in Statistics and Econometrics*. Wiley Series in Probability and Statistics. John Wiley & Sons Ltd., Chichester.
- Mazumder, R., Hastie, T., and Tibshirani, R. (2010). Spectral regularization algorithms for learning large incomplete matrices. *Journal of Machine Learning Research*, 11:2287–2322.
- McLachlan, G. J. and Krishnan, T. (2008). *The EM Algorithm and Extensions*. Wiley Series in Probability and Statistics. Wiley-Interscience [John Wiley & Sons], Hoboken, NJ, second edition.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. (1953). Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092.
- Metropolis, N. and Ulam, S. (1949). The Monte Carlo method. *J. Amer. Statist. Assoc.*, 44:335–341.
- Nocedal, J. and Wright, S. J. (2006). *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer, New York, second edition.
- Novembre, J., Johnson, T., Bryc, K., Kutalik, Z., Boyko, A. R., Auton, A., Indap, A., King, K. S., Bergmann, S., Nelson, M. R., Stephens, M., and Bustamante, C. D. (2008). Genes mirror geography within europe. *Nature*, 456(7218):98–101.
- Price, A. L., Patterson, N. J., Plenge, R. M., Weinblatt, M. E., Shadick, N. A., and Reich, D. (2006). Principal components analysis corrects for stratification in genome-wide association studies. *Nature Genetics*, 38(8):904–909.

- Robert, C. P. and Casella, G. (2004). *Monte Carlo Statistical Methods*. Springer Texts in Statistics. Springer-Verlag, New York, second edition.
- Saad, Y. (2003). *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics, Philadelphia, PA, second edition.
- Saloff-Coste, L. (1997). Lectures on finite Markov chains. In *Lectures on probability theory and statistics (Saint-Flour, 1996)*, volume 1665 of *Lecture Notes in Math.*, pages 301–413. Springer, Berlin.
- Seneta, E. (1996). Markov and the birth of chain dependence theory. *International Statistical Review*, 64:255–263.
- Stein, M., Chen, J., and Anitescu, M. (2012). Difference filter preconditioning for large covariance matrices. *SIAM Journal on Matrix Analysis and Applications*, 33(1):52–72.
- Tanner, M. A. and Wong, W. H. (1987). The calculation of posterior distributions by data augmentation. *J. Amer. Statist. Assoc.*, 82(398):528–550. With discussion and with a reply by the authors.
- Teets, D. and Whitehead, K. (1999). The discovery of Ceres: how Gauss became famous. *Math. Mag.*, 72(2):83–93.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B*, 58(1):267–288.
- Tukey, J. W. (1962). The future of data analysis. *Ann. Math. Statist.*, 33:1–67.
- Vardi, Y., Shepp, L. A., and Kaufman, L. (1985). A statistical model for positron emission tomography. *J. Amer. Statist. Assoc.*, 80(389):8–37. With discussion.