

EECS 498 System Design of a Search Engine

Winter 2026

Lecture 2: Search engine basics

Nicole Hamilton

[https://web.eecs.umich.edu/~nham/
nham@umich.edu](https://web.eecs.umich.edu/~nham/nham@umich.edu)

Agenda

1. Course details
2. MostPositiveSubsequence
3. History of the web
4. Search basics
5. Your project

agenda

1. Course details
2. MostPositiveSubsequence
3. History of the web
4. Search basics
5. Your project

details

1. I've typed your ideas from Friday lab into the slides and posted a PDF to Canvas.
2. We will be doing a survey to pick a mascot for a tee shirt.
3. Your next homework will ask you to create a Gantt chart for a plan to apply to graduate school, due Sunday.
4. I'm hoping everyone can find a team this week. (More speed dating today if we have time.)
5. Questions?

details

1. I've typed your ideas from Friday lab into the slides and posted a PDF to Canvas.
2. We will be doing a survey to pick a mascot for a tee shirt.
3. OH will be every day.
4. your next homework will ask you to create a Gantt chart for a plan to apply to graduate school, due Sunday.
4. I'm hoping everyone can find a team this week. (More speed dating today if we have time.)
5. Questions?

agenda

1. Course details
2. **MostPositiveSubsequence**
3. History of the web
4. Search basics
5. Your project

MostPositiveSubsequence

Write a C++ function that can scan a sequence of N integers in an array A, returning the sum and the left and right indices of the most positive subsequence.

For example, for the sequence

{ -1, 3, 5, 6, -2, -4, 1, 7, -15, 12, 7, -5 }

the best sum = 20, left = 1, right = 10.

MostPositiveSubsequence

Was it ambiguous?

Anything not specified?

MostPositiveSubsequence

Was it ambiguous? Yes.

Anything not specified?

1. What happens if a null sequence is given? (Can you have null sequence?)
2. What if all the numbers are negative or zero?
3. Does it have to be the shortest sequence?
4. What if two sequences are identical?
5. Does performance matter?

what you saw


Rules you were asked to discover:

1. Since an empty argv was treated as a usage error, you were free to assert $n > 0$ in your `MostPositiveSubsequence( )` function.
2. If all the elements are negative, the result is the single least negative element.
3. If two subsequences have identical sums, the shortest, leftmost is preferred.
4. Also, performance and code size does matter. You had to be within 20% of my benchmark and less than 50 LOC.


why shortest first?

Consider this as a signal processing problem, of picking the peak.
Two common characteristics of signals are that they may have harmonics and they may repeat.


What should you choose?



What should you choose?



What should you choose?



naïve solution

```
int NaiveSolution( int a[ ], int n, int &bestL, int &bestR )
{
    assert( n );

    int sum, bestSum, L, R;
    bestSum = a[ 0 ];
    bestL = bestR = 0;

    for ( L = 0; L < n; L++ )
        for ( R = L, sum = 0; R < n; R++ )
        {
            sum += a[ R ];
            if ( sum > bestSum )
            {
                bestL = L;
                bestR = R;
                bestSum = sum;
            }
        }

    return bestSum;
}
```

What makes this naïve?

naïve solution

```
int NaiveSolution( int a[ ], int n, int &bestL, int &bestR )
{
    assert( n );

    int sum, bestSum, L, R;
    bestSum = a[ 0 ];
    bestL = bestR = 0;

    for ( L = 0; L < n; L++ )
        for ( R = L, sum = 0; R < n; R++ )
        {
            sum += a[ R ];
            if ( sum > bestSum )
            {
                bestL = L;
                bestR = R;
                bestSum = sum;
            }
        }

    return bestSum;
}
```

What makes this naïve?

Once a sum ≤ 0 , it can't be useful.

```

int Improved( int a[ ], int n, int &bestL, int &bestR )
{
    assert( n );

    int sum, bestSum, L, R;
    bestSum = a[ 0 ];
    bestL = bestR = 0;

    for ( L = 0; L < n; L++ )
        for ( R = L, sum = 0; R < n; R++ )
        {
            sum += a[ R ];
            if ( sum > bestSum )
            {
                bestL = L;
                bestR = R;
                bestSum = sum;
            }
            if ( sum <= 0 )
                break;
        }

    return bestSum;
}

```

*Stops when the sum ≤ 0 .
Does not find the shortest
sequence.*

```

int Shortest( int a[ ], int n, int &bestL, int &bestR )
{
    assert( n );

    int sum, bestSum, L, R;
    bestSum = a[ 0 ];
    bestL = bestR = 0;

    for ( L = 0; L < n; L++ )
        for ( R = L, sum = 0; R < n; R++ )
        {
            sum += a[ R ];
            if ( sum > bestSum ||
                sum == bestSum && R - L < bestR - bestL )
            {
                bestL = L;
                bestR = R;
                bestSum = sum;
            }
            if ( sum <= 0 )
                break;
        }

    return bestSum;
}

```

Finds the shortest sequence.

```
int Shortest( int a[ ], int n, int &bestL, int &bestR )
{
    assert( n );
```

```
    int sum, bestSum, L, R;
    bestSum = a[ 0 ];
    bestL = bestR = 0;
```

So, what's wrong here?

```
    for ( L = 0; L < n; L++ )
        for ( R = L, sum = 0; R < n; R++ )
        {
            sum += a[ R ];
            if ( sum > bestSum ||
                sum == bestSum && R - L < bestR - bestL )
            {
                bestL = L;
                bestR = R;
                bestSum = sum;
            }
            if ( sum <= 0 )
                break;
        }

    return bestSum;
}
```

```
Student ticks    = 5278229
Benchmark ticks = 260295
Ratio           = 20.2779
Perf limit      = 1.1
Performance is slow
Test has failed
```

```

int Shortest( int a[ ], int n, int &bestL, int &bestR )
{
    assert( n );

    int sum, bestSum, L, R;
    bestSum = a[ 0 ];
    bestL = bestR = 0;

```

Do we really need two nested loops?

```

    for ( L = 0; L < n; L++ )
        for ( R = L, sum = 0; R < n; R++ )
        {
            sum += a[ R ];
            if ( sum > bestSum ||
                sum == bestSum && R - L < bestR - bestL )
            {
                bestL = L;
                bestR = R;
                bestSum = sum;
            }
            if ( sum <= 0 )
                break;
        }

```

Two exits to the for loop.

```

    return bestSum;
}

```

```

int BenchmarkSolution( int a[ ], int n, int &bestL, int &bestR )
{
    assert( n );
    int sum, bestSum, L, R;
    bestSum = a[ 0 ];
    bestL = bestR = 0;

    for ( L = R = sum = 0; R < n; R++ )
    {
        sum += a[ R ];
        if ( sum > bestSum ||
            sum == bestSum && R - L < bestR - bestL )
        {
            bestL = L;
            bestR = R;
            bestSum = sum;
        }
        if ( sum <= 0 )
        {
            sum = 0;
            L = R + 1;
        }
    }
    return bestSum;
}

```

Can this be improved?

Yes, you could move the increment.

Student ticks	= 251726
Benchmark ticks	= 260461
Ratio	= 0.966463
Perf limit	= 1.1
Performance is okay	
Test succeeded	

```
File Edit View Git Project Build Debug Test Analyze Tools Extensions Window Help Search Solution1
Miscellaneous Files (Global Scope)
19
20 #define TC( s ) s, sizeof( s )/sizeof( s[ 0 ] )
21
22 const int a[ ] = { -1, 3, 5, 6, -2, -4, 1, 7, -15, 12, 7, -5 },
23 b[ ] = { -1, 3, 5, 6, -2, -4, 1, 7, -15, 12, 7, -5, -50, 70 },
24 c[ ] = { -1, 3, 5, 6, -2, -4, 1, 7, -15, 12, 7, -5, -50, 70, -69, 6 },
25 d[ ] = { -1, 3, 5, 6, -2, -4, 1, 7, -15, 12, 7, -5, -50, 70, -69, 70 },
26 e[ ] = { -1, 3, 5, 6, -2, 1, 7, -15, 12, 7, -5 },
27 f[ ] = { -1, 3, 5, 6, -2, 1, 7, -15, 12, 7, -5, -50, 29 },
28 g[ ] = { -1, 3, 5, 6, -2, 1, 7, -15, 12, 7, -5, 29 },
29 h[ ] = { -5, -5, -3, -2, -8 },
30 i[ ] = { -5, -5, 0, -1 },
31 j[ ] = { -5, -6, 0, -1 },
32 k[ ] = { -86, -5, -6, -12, -1 },
33 l[ ] = { -86, -5, -6, -12, -1, 2, -1, 5 },
34 m[ ] = { 0, -5, -6, 0, -1 },
35 n[ ] = { 0, -5, -6, 0, -12 },
36 o[ ] = { 1, 2, 0, 3, 0, 0, 0, 2, 1, 3 },
37 p[ ] = { 6, -100, 1, 2, 0, 3, 0, 0, -50, 0, 2, 1, 3 },
38
39 // Naive fails
40 q[ ] = { -5, -5, -3, 0, 0, 5 },
41 r[ ] = { -86, -5, -6, -12, -1, 0, 0, 1, 0 },
42 s[ ] = { 0, -5, -6, 0, 1 },
43 t[ ] = { 1, 2, 0, 3, 0, 0, -50, 0, 2, 1, 3 },
44
45 // Not shortest fails
46 u[ ] = { 5, -100, 1, 2, 0, 3, 0, -25, 6, -12, 0, -50, 0, 2, 1, 3 },
47
48 // Centerpulse
49 v[ ] = { 0, 0, 1, 1, 2, 4, 8, 13, 22, 32, 46, 60, 75, 88, 97, 100,
50 97, 88, 75, 60, 46, 32, 22, 13, 8, 4, 2, 1, 1, 0, 0, 0, 0,
51 0, 0, 1, 4, 14, 35, 72, 120, 164, 181, 164, 120, 72, 35,
52 14, 4, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 1, 2, 4, 8, 13,
53 22, 32, 46, 60, 75, 88, 97, 100, 97, 88, 75, 60, 46, 32,
54 22, 13, 8, 4, 2, 1, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0,
55 0, 0, 0, 0 },
56
57 // CenterpulseDerivative
58 w[ ] = { 0, 0, 0, 1, 1, 2, 4, 6, 8, 11, 13, 15, 15, 13, 9, 3, -3,
59 -9, -13, -15, -15, -13, -11, -8, -6, -4, -2, -1, -1, 0,
60 0, 0, 0, 0, 0, 1, 3, 9, 21, 37, 48, 44, 18, -18, -44,
61 -48, -37, -21, -9, -3, -1, 0, 0, 0, 0, 0, 0, 0, 0, 0,
62 1, 1, 2, 4, 6, 8, 11, 13, 15, 15, 13, 9, 3, -3, -9, -13,
63 -15, -15, -13, -11, -8, -6, -4, -1, -1, -1, 0, 0, 0, 0,
64 0, 0, 0, 0, 0, 0, 0, 0, 0, 0 };
65
66 const TestCase TestSet[ ] =
67 {
68 TC( a ), TC( b ), TC( c ), TC( d ), TC( e ), TC( f ), TC( g ), TC( h ), TC( i ),
69 TC( j ), TC( k ), TC( l ), TC( m ), TC( n ), TC( o ), TC( p ), TC( q ), TC( r ),
70 TC( s ), TC( t ), TC( u ), TC( v ), TC( w )
71 };
72
73 const int TestSetSize = sizeof( TestSet )/sizeof( TestSet[ 0 ] );
74
```

```

int ImprovedBenchmark( int a[ ], int n, int &bestL, int &bestR )
{
    assert( n );
    int sum, bestSum, L, R;
    bestSum = a[ 0 ];
    bestL = bestR = 0;

    for ( L = R = sum = 0; R < n; )
    {
        sum += a[ R ];
        if ( sum > bestSum ||
            sum == bestSum && R - L < bestR - bestL )
        {
            bestL = L;
            bestR = R;
            bestSum = sum;
        }
        R++;
        if ( sum <= 0 )
        {
            sum = 0;
            L = R;
        }
    }
    return bestSum;
}

```

```

Student ticks    = 240777
Benchmark ticks  = 304328
Ratio            = 0.791176
Perf limit       = 1.1
Performance is okay
Test succeeded

```

agenda

1. Course details
2. MostPositiveSubsequence
3. History of the web
4. Search basics
5. Your project

before the web

- 1970s Internal networks, mostly for email. Dial in for time sharing. 300 baud modems.
- 1980s USENET newsgroups. Posts were passed by polling from one machine to the next. 2400 baud modems.
- 1990s Bulletin boards, CompuServe, BIX, many others. 9600 baud modems.

Newsgroups: rec.humor
Path:
gmd.de!xlink.net!sol.ctr.columbia.edu!news.kei.com!ub!acsu.buffalo.edu!ubvms.cc
.buffalo.edu!v140pxgt
From: v140...@ubvms.cc.buffalo.edu (Daniel B Case)
Subject: Canonical List of Light Bulb Jokes
Message-ID: <CEr2F0.521@acsu.buffalo.edu>
News-Software: VAX/VMS VNEWS 1.41
Sender: nn...@acsu.buffalo.edu
Nntp-Posting-Host: ubvmsd.cc.buffalo.edu
Organization: University at Buffalo
Date: Mon, 11 Oct 1993 20:34:00 GMT
Lines: 1021

Sorry for the delay this month, but I had to clear some disk space to post this (In the future, whenever I get there, I will be using Unix so this won't happen-I hope). Anyway, here's the latest version.

:
:

<https://groups.google.com/forum/?hl=en#!original/rec.humor/uRVacQrozNs/Qrmmizp8bUUJ>

Newsgroups: rec.humor

Path:

gmd.de!xlink.net!sol.ctr.columbia.edu!news.kei.com!ub!acsu.buffalo.edu!ubvms.cc
.buffalo.edu!v140pxgt

From: v140...@ubvms.cc.buffalo.edu (Daniel B Case)

Subject: Canonical List of Light Bulb Jokes

Message-ID: <CEr2F0.521@acsu.buffalo.edu>

News-Software: VAX/VMS VNEWS 1.41

Sender: nn...@acsu.buffalo.edu

Nntp-Posting-Host: ubvmsd.cc.buffalo.edu

Organization: University at Buffalo

Date: Mon, 11 Oct 1993 20:34:00 GMT

Lines: 1021

Sorry for the delay this month, but I had to clear some disk space to post this (In the future, whenever I get there, I will be using Unix so this won't happen-I hope). Anyway, here's the latest version.

:
:

<https://groups.google.com/forum/?hl=en#!original/rec.humor/uRVacQrozNs/Qrmimizp8bUUJ>

Newsgroups: rec.humor
Path:
gmd.de!xlink.net!sol.ctr.columbia.edu!news.kei.com!ub!acsu.buffalo.edu!ubvms.cc
.buffalo.edu!v140pxgt
From: v140...@ubvms.cc.buffalo.edu (Daniel B Case)
Subject: Canonical List of Light Bulb Jokes
Message-ID: <CER2F0.521@acsu.buffalo.edu>
News-Software: VAX/VMS VNEWS 1.41
Sender: nn...@acsu.buffalo.edu
Nntp-Posting-Host: ubvmsd.cc.buffalo.edu
Organization: University at Buffalo
Date: Mon, 11 Oct 1993 20:34:00 GMT
Lines: 1021

Sorry for the delay this month, but I had to clear some disk space to post this (In the future, whenever I get there, I will be using Unix so this won't happen-I hope). Anyway, here's the latest version.

:
:

<https://groups.google.com/forum/?hl=en#!original/rec.humor/uRVacQrozNs/Qrmmizp8bUUJ>

the web is born

Aug 6, 1991, the WWW goes live.

1993 Mosaic browser introduced.

1994, WebCrawler and Lycos go live.



Search the web and show for results

Example: "Alien Abduction" UFO Roswell [Search tips](#)



► Choose one of these categories:

[Arts & Literature](#) - [Business](#) - [Chat](#) - [Computers](#) - [Daily News](#)
[Education](#) - [Entertainment](#) - [Government](#) - [Health & Medicine](#)
[Internet](#) - [Kids & Families](#) - [Life & Culture](#) - [Personal Finance](#)
[Recreation](#) - [Reference Desk](#) - [Science](#) - [Sports](#) - [Travel](#)

Get the latest dirt in our new [Gossip](#) section!

Search · [Browse](#) · [Special](#) · [Add URL](#) · [Help](#)


Copyright © 1996 [America Online, Inc.](#)
[Disclaimer](#)



web surfing is born

Through the 90s, individuals and companies post pages with quirky links.

Web surfing begins.



<https://web.archive.org/web/19970601025518/http://www.hamiltonlabs.com:80/>

Links

We've assembled here links to quite a number of interesting sites we've discovered. As you can see, links on this page are *not* all purely business-related. For example, the kids' links are here simply because we have children (as many of you do also) and finding good clean sites you can safely let your kids explore is not always all that easy. If you have some recommendations for additions or if you discover any of these pages have disappeared, please [let us know](#).

Links are broken up into the following categories:

- [Geek Stuff](#)
- [Global Positioning Systems](#)
- [Humor](#)
- [Kids' Stuff](#)
- [News Organizations](#)
- [Search Engines](#)
- [Software Sites](#)
- [Spam and Privacy Issues](#)
- [Stock Market](#)
- [Weather](#)
- [Web Page Creation](#)
- [WinNT Information](#)

(This page was created from our own Internet Explorer Favorites list with [a C shell script](#).)

Geek Stuff:

Almost everything you ever wanted to know is out there somewhere on the internet. There's also a lot of stuff you really didn't care to know. You decide.

Humor:

The title says it all, though we admit that not all the sites we list were actually intended as humorous. Sometimes, that's the whole reason they're so funny.

[\\$95.093](#)

[About The True Mirror](#)

[Application to be Sarang's Girlfriend](#)

[Bill Gates Personal Wealth Clock](#)

[Bill Gates for President 1996!](#)

[Canonical Light Bulb Jokes](#)

[Daily Deaffirmations](#)

[Dave Letterman's Top Ten Lists](#)

[Dead People Server](#)

[Dear Abby](#)

[Dilbert Newsletter](#)

[Geek Site of the Day](#)

[IBM buys Episcopal Church](#)

[Kevin Kelm's Virtual Vomit](#)

[Light Bulb Jokes for Programmers](#)

[Listing of Bondage Stories](#)

[Make Dogs Fast!!](#)

[Microsnot Corporation](#)

[Microsoft Bids to Acquire Catholic Church](#)

[Microsoft Light Bulb Jokes](#)

[Mystical Smoking Head of 'Bob'](#)

[News of the Weird Archives](#)

[Official Outhouses of America Tour](#)

[Other Crackpot Religions!](#)

[Ovi's World of the Bizarre](#)

[Pizza Server!](#)

[Poker at the Dead Eye Saloon](#)

Yahoo!

Mar 2, 1995, Yahoo! goes live but their links are all hand-curated.



[Exploring Mars](#)

The New York Times
Win a **\$20,000 Trip!** [CLICK HERE!](#)

Looking for a
[Car?](#) [Job?](#) [Date?](#)

[options](#)

[Yellow Pages](#) - [People Search](#) - [Maps](#) - [Classifieds](#) - [News](#) - [Stock Quotes](#) - [Sports Scores](#)

- [Arts and Humanities](#)
[Architecture](#), [Photography](#), [Literature](#)...
- [Business and Economy](#) [Xtra!]
[Companies](#), [Investing](#), [Employment](#)...
- [Computers and Internet](#) [Xtra!]
[Internet](#), [WWW](#), [Software](#), [Multimedia](#)...
- [Education](#)
[Universities](#), [K-12](#), [College Entrance](#)...
- [Entertainment](#) [Xtra!]
[Cool Links](#), [Movies](#), [Music](#), [Humor](#)...
- [Government](#)
[Military](#), [Politics](#) [Xtra!], [Law](#), [Taxes](#)...
- [Health](#) [Xtra!]
[Medicine](#), [Drugs](#), [Diseases](#), [Fitness](#)...
- [News and Media](#) [Xtra!]
[Current Events](#), [Magazines](#), [TV](#), [Newspapers](#)...
- [Recreation and Sports](#) [Xtra!]
[Sports](#), [Games](#), [Travel](#), [Autos](#), [Outdoors](#)...
- [Reference](#)
[Libraries](#), [Dictionaries](#), [Phone Numbers](#)...
- [Regional](#)
[Countries](#), [Regions](#), [U.S. States](#)...
- [Science](#)
[CS](#), [Biology](#), [Astronomy](#), [Engineering](#)...
- [Social Science](#)
[Anthropology](#), [Sociology](#), [Economics](#)...
- [Society and Culture](#)
[People](#), [Environment](#), [Religion](#)...

[My Yahoo!](#) - [Yahoo!igans! for Kids](#) - [Beatrice's Web Guide](#) - [Yahoo! Internet Life](#)
[Weekly Picks](#) - [Today's Web Events](#) - [Chat](#) - [Weather Forecasts](#)
[Random Yahoo! Link](#) - [Yahoo! Shop](#)

National Yahoos [Canada](#) - [France](#) - [Germany](#) - [Japan](#) - [U.K. & Ireland](#)

[Yahoo! Africa](#) - [Yahoo! Asia](#) - [Yahoo! Australia](#) - [Yahoo! Brazil](#) - [Yahoo! Canada](#) - [Yahoo! China](#) - [Yahoo! Europe](#) - [Yahoo! France](#) - [Yahoo! Germany](#) - [Yahoo! India](#) - [Yahoo! Japan](#) - [Yahoo! Korea](#) - [Yahoo! Mexico](#) - [Yahoo! New Zealand](#) - [Yahoo! Philippines](#) - [Yahoo! Singapore](#) - [Yahoo! South Africa](#) - [Yahoo! Spain](#) - [Yahoo! Sweden](#) - [Yahoo! Switzerland](#) - [Yahoo! Taiwan](#) - [Yahoo! Thailand](#) - [Yahoo! UK & Ireland](#) - [Yahoo! USA](#) - [Yahoo! Vietnam](#)

Yahoo!

Jerry Yang and David Filo, both PhD students in engineering at Stanford, start Jerry's Guide to the Web.

Jerry hated the name, so it became Yahoo!, a meaningless acronym.

Hand-crafted directory to the web.

Jerry and David categorized 1000 sites/day.

Initially running out of a trailer on campus.

Moved to free hosting on Netscape's machines in early 1995.

"It was a labor love – lots of labor, since no software program could evaluate and categorize sites."

AltaVista

Dec 15, 1995, DEC Research's AltaVista goes live with the first modern search engine architecture.

Used index stream readers invented by Mike Burrows.

AltaVista quickly becomes most popular engine.

Bought by Yahoo! in 2003.



Try your search in: [Shopping](#) • [Images](#) • [Video](#) • [MP3/Audio](#) • [News](#) • [Autos](#) • [Tech](#)

Search for:

[Help](#) | [Customize Settings](#) | [Family Filter is off](#)

any language ▾

[Search](#)

[Search Assistant](#) | [Advanced Search](#)

Shopping: [Compare Prices](#) • [Local Deals & Coupons](#) • [Web Deals & Rebates](#) • [uBid Auction](#)

Tools: [Email](#) • [Translate](#) • [Maps](#) • [Directions](#) • [Yellow Pages](#) • [People Finder](#) • [Find A Date](#)
[Find Downloads](#) • [Text-Only Search](#) • [Weight Calculator](#) • [Find A Job](#) • [Find A Home](#) • [More...](#)

News: [Supreme Court Upholds Use of Force in Guarding...](#) • [More News...](#)

[Web Site Hosting](#) • [Insurance Quotes](#) • [Radio Pet Fence](#) • [Get Cash Back](#) • [Buy A Computer](#)
[Vacation Plans](#) • [Online Casinos & Gambling](#) • [DVD Players](#) • [NBA Tickets](#) • [Win Free Travel](#)

Arts & Entertainment

[Culture](#), [Celebrities](#), [Movies...](#)

Music

[Artists](#), [Genres](#), [MP3...](#)

Business Center

[Internet Search Services](#)

[AltaVista Enterprise Software](#)

[Download a trial version now](#)

[Submit A Site](#)

[List Your Products](#)

[Advertise With Us](#)

[Access Files From Anywhere](#)

Autos

[Buy & Sell](#), [Guides](#), [Repair...](#)

People & Chat

[Chat](#), [Email](#), [Personals...](#)

Computing

[Hardware](#), [Internet](#), [Software...](#)

Personal

[Family](#), [Intimacy](#), [Home...](#)

Games

[Gambling](#), [Role Playing](#), [Video...](#)

Travel

[Activities](#), [Destinations](#), [Trips...](#)

Health & Fitness

[Conditions](#), [Medicine](#), [Insurance...](#)

Shopping

[Coupons](#), [Deals](#), [Wireless...](#)

Google

Sep 4, 1998, Google goes live with PageRank, named after Larry Page, coinvented by Sergey Brin.

Introduced a much simpler page where they promised that all their results were algorithmic, none paid and that *they would do no evil*.

Quickly became the dominant engine.



Search 1,326,920,000 web pages

[Advanced Search](#)
[Preferences](#)


Google Search I'm Feeling Lucky

Google Web Directory
the web organized by topic

[Cool Jobs](#) - [Add Google to Your Site](#) - [Advertise with Us](#) - [Google in Your Language](#) - [All About Google](#)


©2001 Google

Why did this become dominant?



Bing

Jan 2005, Microsoft's own engine goes live, initially as MSN Search.



msn. Search

[Web](#) [News](#) [Images](#) [Desktop](#) [Encarta](#)

[Search](#) [Near Me](#)

[+Search Builder](#) [Settings](#) [Help](#) [Español](#)

It's a lot easier to do something if you know it's possible.

In business, this called either an *imitator* or *follower* strategy.

In engineering, if you know something's possible, you know not to give up.


agenda

1. Course details
2. MostPositiveSubsequence
3. History of the web
4. Search basics
5. Your project


the basic parts to a search engine

1. HTML parser.
2. Crawler.
3. Index.
4. Constraint solver.
5. Query language.
6. Ranker.
7. Front end.


the index build side




the index build side




the query serve side



the query serve side




HTML Parser



Extract the content from a HTML file as a series of tokens in the title and the body of the document and a set of links with anchor text to other documents.

HTML Parser




Historically, many teams have found that getting the HTML parser to work is the most time-consuming part of the project.

There is so much bad, broken HTML out there on the web that many teams struggle to keep their parsers from crashing.

This is a big part of the *surface area* of your project.

HTML Parser



Common pitfall: Overengineering and overcomplexity.

Think carefully: What does it *need* to do?

Does it *need* to parse the entire document structure?

You have a homework that should help point you in the right direction.

HTML

```
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
    "http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
<html xmlns="http://www.w3.org/1999/xhtml">

<head>
    <meta content="text/html; charset=utf-8" http-equiv="Content-Type" />


    <title>Page title shown in the browser tab</title>
    <script>Lots of gibberish</script>

    <link href="MyStyles.css" rel="stylesheet" type="text/css" />
    <link href="https://mydomain/favicon.ico" rel="shortcut icon" />

    <meta name="DC.Rights" content="Copyright 2018 my name"/>
    <meta name="description" content="Short abstract."/>
    <meta name="keywords" content="arbitrary list of words"/>
</head>


<body>
:
</body>
</html>
```

crawler



1. Manage a frontier of new links to be crawled.
2. Decide what will or will not be crawled and in what order.
3. Keep track of what's already been crawled.
4. Read pages over HTTP and HTTPS.
5. Obey robots.txt files.
6. Deal with redirects.

crawler




All of this has to be highly multithreaded so you don't wait on slow sites, instead overlapping them.

You will also want to spread it across multiple machines and will need to decide how to divide it up and pass information.

This also is part of your *surface area*.

crawler




Typically maintains pool of worker processes or threads to read and parse webpages.

Each worker retrieves the file and queues it for the HTML parser which creates an object.

Links go into the frontier, perhaps with anchor.

Words go into the index.


crawler



It's very easy to DOS a site with thousands of threads on a bunch of AWS machines with gigabit connections.

So, we'll talk about politeness and putting **your** contact info in your *User-Agent*: field so complaints go to you, not me.

index



A merged *inverted word index* of all the documents that have been crawled, allowing you to report all the documents and individual locations (postings) where any given word was found.

Due to the size, the posting lists will have to be on disk but you'll map them right into your process virtual memory address space.

The inverted word index operates like the index at the back of a book.

An index in a book lists all pages where a word or topic appears.

It's called inverted because we think of the original text as the non-inverted form.

INDEX AND GLOSSARY

*Some Men pretend to understand a Book
by scouting thro' the Index:
as if a Traveller should go about to describe a Palace
when he had seen nothing but the Privy.*

—JONATHAN SWIFT
(*Mechanical Operation of the Spirit*, 1704)

When an index entry refers to a page containing a relevant exercise, see also the *answer* to that exercise for further information; an answer page is not indexed here unless it refers to a topic not included in the statement of the exercise.

- A-register of MIX, 122.
- A-1 compiler, 458.
- Aardenne-Ehrenfest, Taniana van, 375, 578.
- Aarons, Roger M., 522.
- Abel, Niels Henrik, 56.
- binomial formula generalized, 56, 70, 72, 398.
- limit theorem, 94.
- Abramowitz, Milton, 66, 92, 615.
- ACE computer, Pilot, 226.
- Adams, Charles William, 226.
- ADD, 127, 128, 204.
- Add to list: *see* Insertion.
- Addition of polynomials, 273–276, 355–359, 361.
- Address: A number used to identify a position in memory.
- field of MIXAL line, 123, 141, 147, 151, 152.
- of node, 229–230.
- portion of MIX instruction, 123.
- Address transfer operators of MIX, 129, 206–207.
- Adjacent vertices of a graph, 362.
- Agenda, 285, 293, *see* Priority queue.
- Ahrens, Wilhelm Ernst Martin Georg, 159.
- al-Khowārizmī, Abu Ja'far Mohammed ibn Mūsā, 1, 78.
- Alanen, Jack David, xiii.
- ALF (alphabetic data), 148, 149, 151.
- Algebraic formulas, manipulation of, 335–347, 461.
- differentiation, 337–346, 359, 458.
- representation as trees, 312, 335–336, 458.
- simplification of, 339, 346.
- Algorithms, origin of word, 1–2.
- Algorithms, 1–9.
- analysis of, vii, 7, 94–104, 166–169, 175, 246–247, 249–250, 265, 276, 323–324, 380–381, 445–446.
- communication of, 16.
- effective, 6, 8, 9.
- equivalence between, 466.
- form of in this book, 2–4.
- hardware-oriented, 26, 249, 600.
- how to read, 4, 16.
- proof of, 14–20, 318–319, 420, 566.
- properties of, 4–6, 9.
- random paths in, 380–381.
- set theoretical definition, 8–9.
- theory of, 7, 9.
- Allocation of tables, *see* Dynamic storage allocation, Linked allocation, Representation, Sequential allocation.
- Along order, 459.
- Alphameric character: A letter, digit, or special character symbol.
- codes for MIX, 132, 134, 136–137.
- AMM: *American Mathematical Monthly*, the official journal of the Mathematical Association of America, Inc.
- Analysis of algorithms, vii, 7, 94–104, 166–169, 175, 246–247, 249–250, 265, 276, 323–324, 380–381, 445–446.
- Analytical Engine, 1, 225.
- Ancestor, in a tree structure, 309.
- André, Antoine Désiré, 531.
- Anticipated input, 212, *see* Buffering.
- Antisymmetric relation, 258.
- Apex of tree, 307.
- Apostol, Tom Mike, 28.
- Arborescence, 362, *see* Oriented trees.
- Arc in a directed graph, 371.
- Are-digraph, 379.
- Area of memory, 435.
- Arguments of subroutines, 183, 185.
- Arithmetic: Addition, subtraction, multiplication, and division, vii.
- fixed-point, 154–157.
- floating-point, 127, 304.
- operators of MIX, 127–128, 135, 204.
- polynomial, 272–277, 355–359, 361.
- scaled decimal, 156–157.
- Arithmetic expressions, *see* Algebraic formulas.
- Arithmetic progression, sum of, 11, 13, 31, 55.

If you want to know the list of pages where some combination of topics appears, all you need to do is find the intersections of the lists.

INDEX AND GLOSSARY

*Some Men pretend to understand a Book
by scouting thro' the Index:
as if a Traveller should go about to describe a Palace
when he had seen nothing but the Privy.*

—JONATHAN SWIFT
(*Mechanical Operation of the Spirit*, 1704)

When an index entry refers to a page containing a relevant exercise, see also the *answer* to that exercise for further information; an answer page is not indexed here unless it refers to a topic not included in the statement of the exercise.

- A-register of MIX, 122.
- A-1 compiler, 458.
- Aardenne-Ehrenfest, Taniana van, 375, 578.
- Aarons, Roger M., 522.
- Abel, Niels Henrik, 56.
 - binomial formula generalized, 56, 70, 72, 398.
 - limit theorem, 94.
- Abramowitz, Milton, 66, 92, 615.
- ACE computer, Pilot, 226.
- Adams, Charles William, 226.
- ADD, 127, 128, 204.
- Add to list: *see* Insertion.
- Addition of polynomials, 273–276, 355–359, 361.
- Address: A number used to identify a position in memory.
 - field of MIXAL line, 123, 141, 147, 151, 152.
 - of node, 229–230.
 - portion of MIX instruction, 123.
- Address transfer operators of MIX, 129, 206–207.
- Adjacent vertices of a graph, 362.
- Agenda, 285, 293, *see* Priority queue.
- Ahrens, Wilhelm Ernst Martin Georg, 159.
- al-Khowārizmī, Abu Ja'far Mohammed ibn Mūsā, 1, 78.
- Alanen, Jack David, xiii.
- ALF (alphabetic data), 148, 149, 151.
- Algebraic formulas, manipulation of,
 - 335–347, 461.
 - differentiation, 337–346, 359, 458.
 - representation as trees, 312, 335–336, 458.
 - simplification of, 339, 346.
- Algorithms, 1–9.
 - analysis of, vii, 7, 94–104, 166–169, 175, 246–247, 249–250, 265, 276, 323–324, 380–381, 445–446.
 - communication of, 16.
 - effective, 6, 8, 9.
 - equivalence between, 466.
 - form of in this book, 2–4.
 - hardware-oriented, 26, 249, 600.
 - how to read, 4, 16.
 - proof of, 14–20, 318–319, 420, 566.
 - properties of, 4–6, 9.
 - random paths in, 380–381.
 - set theoretical definition, 8–9.
 - theory of, 7, 9.
- Allocation of tables, *see* Dynamic storage allocation, Linked allocation, Representation, Sequential allocation.
- Along order, 459.
- Alphameric character: A letter, digit, or special character symbol.
 - codes for MIX, 132, 134, 136–137.
- AMM: *American Mathematical Monthly*, the official journal of the Mathematical Association of America, Inc.
- Analysis of algorithms, vii, 7, 94–104, 166–169, 175, 246–247, 249–250, 265, 276, 323–324, 380–381, 445–446.
- Analytical Engine, 1, 225.
- Ancestor, in a tree structure, 309.
- André, Antoine Désiré, 531.
- Anticipated input, 212, *see* Buffering.
- Antisymmetric relation, 258.
- Apex of tree, 307.
- Apostol, Tom Mike, 28.
- Arborescence, 362, *see* Oriented trees.
- Arc in a directed graph, 371.
- Are-digraph, 379.
- Area of memory, 435.
- Arguments of subroutines, 183, 185.
- Arithmetic: Addition, subtraction, multiplication, and division, vii.
 - fixed-point, 154–157.
 - floating-point, 127, 304.
 - operators of MIX, 127–128, 135, 204.
 - polynomial, 272–277, 355–359, 361.
 - scaled decimal, 156–157.
- Arithmetic expressions, *see* Algebraic formulas.
- Arithmetic progression, sum of, 11, 13, 31, 55.

An inverted word list is a computer scientist's term for an older idea.

INDEX AND GLOSSARY

*Some Men pretend to understand a Book
by scouting thro' the Index:
as if a Traveller should go about to describe a Palace
when he had seen nothing but the Privy.*


—JONATHAN SWIFT
(*Mechanical Operation of the Spirit*, 1704)

When an index entry refers to a page containing a relevant exercise, see also the *answer* to that exercise for further information; an answer page is not indexed here unless it refers to a topic not included in the statement of the exercise.

- A-register of MIX, 122.
- A-1 compiler, 458.
- Aardenne-Ehrenfest, Taniana van, 375, 578.
- Aarons, Roger M., 522.
- Abel, Niels Henrik, 56.
- binomial formula generalized, 56, 70, 72, 398.
- limit theorem, 94.
- Abramowitz, Milton, 66, 92, 615.
- ACE computer, Pilot, 226.
- Adams, Charles William, 226.
- ADD, 127, 128, 204.
- Add to list: *see* Insertion.
- Addition of polynomials, 273–276, 355–359, 361.
- Address: A number used to identify a position in memory.
- field of MIXAL line, 123, 141, 147, 151, 152.
- of node, 229–230.
- portion of MIX instruction, 123.
- Address transfer operators of MIX, 129, 206–207.
- Adjacent vertices of a graph, 362.
- Agenda, 285, 293, *see* Priority queue.
- Ahrens, Wilhelm Ernst Martin Georg, 159.
- al-Khowārizmī, Abu Ja'far Mohammed ibn Mūsā, 1, 78.
- Alanen, Jack David, xiii.
- ALF (alphabetic data), 148, 149, 151.
- Algebraic formulas, manipulation of, 335–347, 461.
- differentiation, 337–346, 359, 458.
- representation as trees, 312, 335–336, 458.
- simplification of, 339, 346.
- Algorithms, origin of word, 1–2.
- analysis of, vii, 7, 94–104, 166–169, 175, 246–247, 249–250, 265, 276, 323–324, 380–381, 445–446.
- communication of, 16.
- effective, 6, 8, 9.
- equivalence between, 466.
- form of in this book, 2–4.
- hardware-oriented, 26, 249, 600.
- how to read, 4, 16.
- proof of, 14–20, 318–319, 420, 566.
- properties of, 4–6, 9.
- random paths in, 380–381.
- set theoretical definition, 8–9.
- theory of, 7, 9.
- Allocation of tables, *see* Dynamic storage allocation, Linked allocation, Representation, Sequential allocation.
- Along order, 459.
- Alphameric character: A letter, digit, or special character symbol.
- codes for MIX, 132, 134, 136–137.
- AMM: *American Mathematical Monthly*, the official journal of the Mathematical Association of America, Inc.
- Analysis of algorithms, vii, 7, 94–104, 166–169, 175, 246–247, 249–250, 265, 276, 323–324, 380–381, 445–446.
- Analytical Engine, 1, 225.
- Ancestor, in a tree structure, 309.
- André, Antoine Désiré, 531.
- Anticipated input, 212, *see* Buffering.
- Antisymmetric relation, 258.
- Apex of tree, 307.
- Apostol, Tom Mike, 28.
- Arborescence, 362, *see* Oriented trees.
- Arc in a directed graph, 371.
- Are-digraph, 379.
- Area of memory, 435.
- Arguments of subroutines, 183, 185.
- Arithmetic: Addition, subtraction, multiplication, and division, vii.
- fixed-point, 154–157.
- floating-point, 127, 304.
- operators of MIX, 127–128, 135, 204.
- polynomial, 272–277, 355–359, 361.
- scaled decimal, 156–157.
- Arithmetic expressions, *see* Algebraic formulas.
- Arithmetic progression, sum of, 11, 13, 31, 55.


Historically, concordances have been used to search important works.

A concordance is an alphabetical list of the words used in a book or body of work, listing every instance with its immediate context.




The first Bible concordance was compiled for the Vulgate Bible by Hugh of St Cher (d.1262), who employed 500 friars to assist him.

The audience was priests preparing sermons.



Here's how an online Bible concordance works.

Notice how it gives the context. For the Bible, that's the whole verse.




<https://thekingsbible.com/Search/SearchResults>

Concordances have also been created for the Vedas, Qur'an, the works of Shakespeare and other authors.

Before computers, they had to be constructed by hand. So, they were only done for important works.

This one is 1000 pages.



Here's how an online
Shakespeare
concordance.

The screenshot shows a web browser window with the URL <https://www.opensourceshakespeare.org/concordance/>. The page title is "Concordance of Shakespeare's complete works". The header includes the "OpenSourceShakespeare" logo and a navigation menu with links: "Plays", "Sonnets", "Poems", "Concordance", "Advanced Search", and "About OSS". A search bar labeled "SEARCH TEXTS" is located in the top right. Below the header, there is a section titled "Concordance of Shakespeare's complete works" with a small red logo. A search form is present with the label "Find a word form" and a text input field. Below the input field are three radio button options: "Exact spelling" (selected), "First part of a word form", and "Any part of a word form". To the left of the main content area is a vertical list of letters from A to L, each followed by a count in parentheses: A (1,569), B (1,847), C (2,625), D (1,754), E (1,067), F (1,459), G (915), H (1,190), I (869), J (252), K (200), and L (1,082). To the right of this list is a paragraph of text explaining the search functionality and providing instructions on how to use the concordance.

Concordance of Shakespeare's complete works

Find a word form

☒ Exact spelling
☐ First part of a word form
☐ Any part of a word form

A (1,569) **Choose a letter** to the left to see all the word forms beginning with that letter. The numbers next to the letters indicate how many word forms begin with that letter.

B (1,847) **Do you have to look through hundreds of words to find the word form you're seeking?** Of course not! The OSS management values your time. Enter the exact word form you want in the box above to search for it.

C (2,625)

D (1,754) **What is a word form?** A word like "play" can take several forms, like *plays*, *playing*, and *played*. Each word form is listed in the database, along with the number of times the word form occurs in the text collection.

E (1,067)

F (1,459)

G (915)

H (1,190)

I (869)

J (252)

K (200)

L (1,082)

<https://www.opensourceshakespeare.org/concordance/>


If I search for stars, I can find all the lines that mention stars and who said it.

The screenshot shows a web browser displaying the OpenSourceShakespeare search results page. The browser's address bar shows the URL <https://www.opensourceshakespeare.org/search/se...>. The page has a navigation bar with links to [Plays](#), [Sonnets](#), [Poems](#), [Concordance](#), [Advanced Search](#), and [About OSS](#). The main heading is "Search results" with a subtext "1-20 of 75 total". The keyword "stars" is displayed in a red box. Below the keyword are three buttons: "Refine your search", "Print-friendly", and "New search". A message states: "For an explanation of each column, tap or hover over the column's title." Below this is a pagination control showing "1" (selected), "2", "3", "4", and "NEXT". The search results are presented in a table with four columns: "#", "Work", "Character", and "Line Text".


#	Work	Character	Line Text
1	All's Well That Ends Well [I, 1]	Helena	182 That wishing well had not a body in't, Which might be felt; that we, the poorer born, Whose baser stars do shut us up in wishes, Might with effects of them follow our friends, And show what we alone must think, which never Return us thanks.
2	All's Well That Ends Well [I, 3]	Helena	573 There's something in't, More than my father's skill, which was the greatest Of his profession, that his good receipt Shall for my legacy be sanctified By the luckiest stars in heaven: and, would your honour But give me leave to try success, I'd venture The well-lost life of mine on his grace's cure By such a day and hour.
3	All's Well That Ends Well [II, 5]	Helena	1341 And ever shall With true observance seek to eke out that Wherein toward me my homely stars have fail'd To equal my great fortune.
4	All's Well That Ends Well [V, 3]	King of France	2945 You boggle shrewdly, every feather stars you. Is this the man you speak of?
5	Antony and Cleopatra [III, 2]	Lepidus	1675 Let all the number of the stars give light To thy fair way!
6	Antony and Cleopatra	Antony	2370 Moon and stars Whip him. Were't twenty of the greatest tributaries

<https://www.opensourceshakespeare.org/concordance/>


In your search engine you'll create an index as file containing the list of URLs contained in the index and the inverted word index.




The inverted word index




A posting list



An individual post




constraint solver



Given an inverted word index and a constraint, e.g., a list of words that must appear together or as a phrase in a document, find the list of matching documents.

constraint solver



Index Stream Readers (ISRs) are used to scan the index.

Finds the next occurrence of the desired token or combination of child ISRs representing individual words, document ends, ORing or ANDing of terms, Phrases, and Nots (terms that should not occur in the document.)

Consider these posting lists

quick	10	27	105	513	518	520
brown	28	50	62	70	514	790
fox	87	106	515	550	1200	
#DocEnd	112	570	1006	1704		

To read and merge these lists, we need to move from one entry to the next.

We'll do that with an ISR (index stream reader).

The ISR for each token has to be able to report its current location and any attributes, and it needs Next() and Seek() functions.

OR'ing streams

quick	10	27	105	513	518	520
brown	28	50	62	70	514	790
fox	87	106	515	550	1200	
#DocEnd	112	570	1006	1704		

quick | fox 10 27 87 105 106 513 515 518 520

An OR ISR simply merges the streams.

No need to pay attention to document boundaries. Each post is in whichever posting list and whatever document it happens to be.

AND'ing streams

quick	10	27	105	513	518	520
brown	28	50	62	70	514	790
fox	87	106	515	550	1200	
#DocEnd	112	570	1006	1704		

quick fox ?

AND'ing streams


Easier to consider if we show the document boundaries.

quick	10	27	105		513	518	520					
brown	28	50	62	70	514				790			
fox	87	106			515	550					1200	
#DocEnd					112			570	1006			1704
quick fox												

?

AND'ing of terms should find occurrences of all the terms within a single document.


query compiler




Compile a query as typed, e.g., with quotes around phrases, into a structure the constraint solver understands.

The query language and the ISRs should be recursive

"apollo moon landing" | (apple banana)




ranker



Use the compiled query to retrieve matching pages from the index, then score them by a statistical, heuristic or other method and pick the ten best.

front end



Report the results back,
possibly including snippets
with a simple HTTP server.

Result quality

To be useful as a search engine, it must produce good results quickly.

result quality

Traditional measures:

Precision The fraction of the results that are relevant.

Recall The fraction of all relevant results returned.

result quality

And because it matters that the top results be the best:

Relevance An estimate of the overall quality of the page as an result for this query that can be used for *ranking* (ordering) the results.

kinds of rank

Dynamic rank

How good a match this page is to this specific query.

Static rank

How good the page is, knowing nothing about the query.

dynamic rank

Estimate the quality of match between the page and the query.

Traditional methods in information retrieval considered the page as a *bag of words* and ostensibly measured the *mathematical* or *statistical similarity*.

For example, one method claimed to consider the *document and the query as vectors* and to calculate the *angle between them*.

bag of words

A scoring method that considers only that the search words occur in the document, not where or in what relation to each other.

tf-idf

Term-frequency, inverse document frequency.

Easily the most famous bag of words technique. The more occurrences of a rare word, the better.

bag of words

Here's a sample NY Times page from Jan 25, 2021.

The New York Times - Breaking News x +

File | C:/Users/hamil/Google%20Drive/eecs440/ExampleCode/HtmlParserSolution/nytimes.html

Seattle PI WSJ WSJ NY Times WaPost Canvas Course Info Center Crappy Scheduler BofA Wikipedia GSI/IA Applications Student evals Patents Other bookmarks

U.S. INTERNATIONAL CANADA ESPAÑOL 中文

The New York Times

Monday, January 25, 2021
Today's Paper

29°F 31° 26°
Nasdaq -0.02% ↓

SUBSCRIBE FOR \$1/WEEK LOG IN

World U.S. Politics N.Y. Business Opinion Tech Science Health Sports Arts Books Style Food Travel Magazine T Magazine Real Estate Video

The Daily Listen to 'The Daily' Impeached, again.

The Book Review Podcast Charles Yu on his National Book Award-winning novel, "Interior Chinatown."

The Argument Listen to 'The Argument' Jane Coaston loves Section 230. Got a problem with that?

LIVE

Fractures in G.O.P. Widen After Impeachment Vote

- Some House Republicans have called on Representative Liz Cheney to resign from her leadership post after she voted to impeach President Trump.
- Ms. Cheney was one of 10 Republicans to break with the party on the vote, setting up a messy internal feud that could define its future.




'Suspicious' Visits to Capitol Eyed as Dragnet Snares 2 Police Officers

Democratic lawmakers demanded answers about visitors to the Capitol a day before the riot.

<https://www.nytimes.com/2021/01/08/books/review/podcast-charles-yu-interior-chinatown-david-brown-henry-adams-last-american-aristocrat.html>

bag of words


Here it is stripped of HTML and CSS but the text remains.



The screenshot shows a web browser window with the address bar displaying a file path: `C:/Users/hamil/Google%20Drive/eecs440/ExampleCode/HtmlParserSolution/Simplepage.html`. The browser's address bar also shows the page title "The New York Times - Breaking News". The page content is a list of links and text, including "Continue reading the main story", "Sections", "SEARCH", "Skip to content", "Skip to site index", "U.S. International Canada Español 中文", "Today's Paper", "World U.S. Politics N.Y. Business Opinion Tech Science Health Sports Arts Books Style Food Travel Magazine T Magazine Real Estate Video World U.S. Politics N.Y. Business Opinion Tech Science Health Sports Arts Books Style Food Travel Magazine T Magazine Real Estate Video Listen to 'The Daily' Impeached, again. The Book Review Podcast Charles Yu on his National Book Award-winning novel, 'Interior Chinatown.' Listen to 'The Argument' Jane Coaston loves Section 230. Got a problem with that? Live Fractures in G.O.P. Widen After Impeachment Vote". The text continues with a paragraph about Representative Liz Cheney and the impeachment of President Trump, followed by a list of links and text including "Some House Republicans have called on Representative Liz Cheney to resign from her leadership post after she voted to impeach President Trump. Ms. Cheney was one of 10 Republicans to break with the party on the vote, setting up a messy internal feud that could define its future. President-elect Joe Biden has stayed largely removed from the fray of the impeachment and kept his focus on the pandemic and the economy. Here's the latest. T.J. Kirkpatrick for The New York Times Erin Schaff/The New York Times T.J. Kirkpatrick for The New York Times T.J. Kirkpatrick for The New York Times Jason Andrew for The New York Times Erin Schaff/The New York Times Rare Pentagon Mission: Armed Troops in Washington for Inauguration The troops will be the National Guard and not active-duty. But officials know such events are fraught with risk. 'Suspicious' Visits to Capitol Eyed as Dragnet Snares 2 Police Officers Democratic lawmakers demanded answers about visitors to the Capitol a day before the riot. The Capitol Police Had One Mission. Now the Force Is in Crisis. The department's performance during last week's rampage has cast a harsh light on an agency unaccustomed to public scrutiny. A man who carried a Confederate flag into the Capitol has been arrested. The F.A.A. ordered airlines to crack down on unruly passengers after an increase in threatening or violent behavior. Biden Set to Call for Big Spending on Stimulus and Vaccines President-elect Joe Biden is expected to lay out plans today for an initial effort to fight the virus and a subsequent one to address the economy. LIVE Sharp Rise in Jobless Claims Underscores Pandemic's Economic Toll A total of 1.15 million workers filed initial claims for state unemployment during the first full week of 2021. Here's the latest on the economy. live Jaime Harrison Said to Be Pick for Next D.N.C. Chair Mr. Harrison, a veteran of South Carolina politics, was the Democratic nominee against Senator Lindsey Graham and shattered fund-raising records in the race. Though Jaime Harrison lost the Senate race in South Carolina, he developed a broad bench of support across the Democratic party. Cameron Pollack for The New York Times A Republican Lawmaker for Whom the Spectacle Is the Point Lauren Boebert, a freshman in the House, evaded metal detectors and ignored police officers before denouncing the impeachment vote. Southwest Surge Helps Drive Record Covid-19 Death Toll in U.S. As the national death toll nears 400,000, weekly deaths in some Arizona and California counties have reached new highs. A coalition of tech companies, health organizations and nonprofits is developing digital vaccination credentials. Here's the latest on the pandemic. live Two Members of W.H.O. Team Are Denied Entry to China A team of experts arrived in Wuhan, but Beijing has repeatedly impeded the global inquiry into the emergence of the coronavirus. The Maddening Red Tape Facing Older People Who Want the Vaccine A million New York City residents who are over 65 are now eligible to be vaccinated. But try making an appointment. Where 2020's Record Heat Was Felt the Most New data shows that 2020 has effectively tied the hottest year on record, making the last seven years the warmest since the beginning of modern record-keeping. The Great Read The Plucky Newsletter That Isn't Afraid to Name Names At 80, the defiant editor of The Buzz "had it with this whole secrecy thing." So she started a newsletter and slipped it under everyone's door. They Can't Leave the Bay Area Fast Enough As a tech era draws to an end, more workers and companies are packing up. What comes next? Keyan Karimi left his \$4,000-a-month one-bedroom in San Francisco for a new opportunity in Savannah, Ga. Stephen B. Morton for The New York Times Opinion Stephen I. Vladeck Why Trump Can Be Convicted Even as an Ex-President He is the poster child for why such accountability is not just constitutionally permissible but necessary. Linda Greenhouse The Supreme Court After Trump Now that his presidency has disintegrated into mayhem, how are the conservative justices processing his tenure? Gail Collins Trump Swings and Misses Stuart A. Thompson and Charlie Warzel They Used To Post Selfies. Now They're Trying to Reverse the Election. Kara Swisher Food Delivery Is Keeping Uber Alive. Will It Kill Restaurants? Suzanne Nossel Yes, Impeach Trump. But Let's Talk About the Meaning of 'Incitement.' Join Kara Swisher to discuss her interview with Uber's C.E.O. Nicholas Kristof Impeachment Is a Call for 'Moral Clarity' Charles M. Blow What Unity? Looking Ahead to Trump's Trial in the Senate Farhad Manjoo The Capitol Was Just the Start Editors' Picks Devin Oктar Yalkin for The New York Times Mike D Is Cleaning Out the Family Attic A collection of furniture, books and memorabilia that belonged to the rapper's mother will go on view and for sale at Sotheby's this month. Seeking a Sunny Manhattan Studio for \$2,000

bag of words

Here it is with the words in the title and body sorted. Tf-idf can't tell the difference.



at Microsoft

At Microsoft, we separated the occurrences into four metastreams, in order of importance.

1. Anchor text.
2. URL.
3. Title.
4. Body text.

heuristics

1. Pick the rarest word in the query, then iterate over its occurrences on the page.
2. At each place, move the pointers to the other words to as close to the correct relative position as possible.
3. Score that that set, called a *span*, counting the number of spans, the number of exact phrases, in order, close together, etc.
4. Score the counts, weighting exact phrases the title as more important than other matches.

heuristics

Altogether, had about 700 heuristics in a linear combination.
Tuned them using gradient descent and approximately 150K
labeled pages and queries.

what we found

Finding the search words in the right order and close together was more important than simply finding lots of them.

Added tf-idf because under pressure that of course we should consider it. Removed it because it added nothing to our relevance but slowed our crawl.

It mattered a lot where we found the matching terms.

Anchor > URL > Title > Body text

static rank


Some pages are just better than others before you know anything about the query.

static rank

1. Some domains are better than others, e.g., .gov or .edu over .biz.
2. Short URLs are better.
3. Short titles are probably better.
4. Some pages may be obvious spam.
5. Some pages may have lots of other pages pointing to them, e.g., PageRank.

PageRank

The basic idea: The more and better links to a page, the more likely it should rank higher.



PageRank

It obviously did work and they got better results.

It also gave halo of special legitimacy to their results, that they were scientific and unbiased.

at Microsoft

We gamely expected our version of PageRank to represent about half the overall rank value, largely based on the hype around it.

Turned out it was very expensive to calculate and represented only a small part of the final rank score.

Agenda

1. Course details.
2. HW1 MostPositiveSubsequence().
3. History of the web.
4. Search basics.
5. Your project.

your project

I've broken it into 8 levels of functionality:

- 0 Basic plan for your project.
- 1 Parse text files into a hash table.
- 2 Build a crawler.
- 3 Build a reverse word index.
- 4 Create a user interface.
- 5 Build a constraint solver and query parser.
- 6 Build a ranker.
- 7 Advanced functionality.

next

Overall course plan is:

1. Cover the OS topics you'll need.
2. Get you crawling ASAP.
3. Work through each piece in order and in time for you to get it done.

Next time: Project planning.

Who needs a team?