Image & Video Compression

Deepak Kumar

Image Compression

- Grayscale raster images
  - 8-bits/pixel – 256 shades of gray
- Color Raster images
  - 24 bits/pixel (R, G, B)
    256x256x256 = 16,777,216 colors
- A 1080p image for HDTV requires 1920 x 1080 = 2,073,600 pixels
  - A grayscale image of this size will require 2,073,600 bytes (1 byte/pixel)
  - A color image of this size will require 1920 x 1080 x 3 = 6,220,800 bytes (3 bytes/pixel)
- Several compression schemes exist for compressing image files depending on the type of image and perceptual quality desired.

Lossless Image Compression

- No data is lost from the original image
- Employs LZW or LZ77 in GIF, PNG, and TIFF file formats
- JPEG-LS is a lossless JPEG format. Slightly better compression ratios.
- Compression ratios are typically ~2:1 for natural imagery but can be much larger for document images.

Lossy Image Compression

- Data is lost from the original image to achieve higher compression
- Exploits perceptual redundancy in image data

Lossy Image Compression

- Data is lost from the original image to achieve higher compression
- Exploits perceptual redundancy in image data

e.g. spectral response of human eye

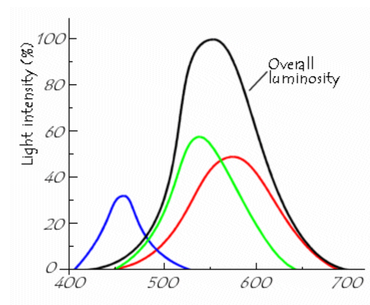

JPEG Compression: General Framework

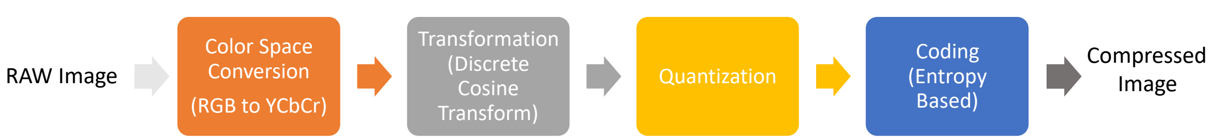

JPEG Compression: General Framework

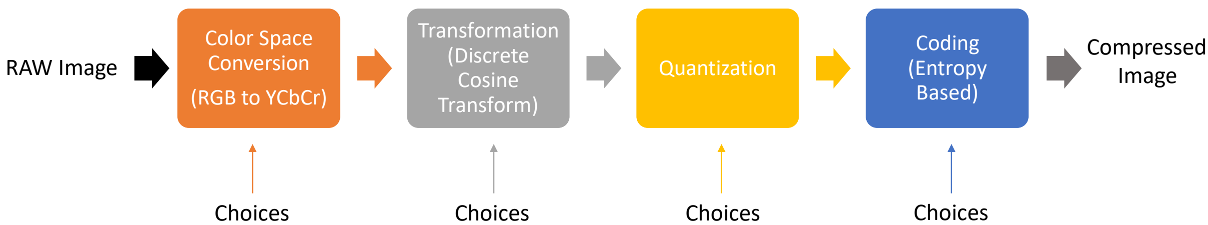

10/7/2019 7

Lossy Image Compression

- Data is lost from the original image to achieve higher compression
- Exploits perceptual redundancy in image data

  e.g. spectral response of human eye

  BTW, to convert a color image into Grayscale use

  gray = 0.2989*r + 0.5870*g + 0.1140*b

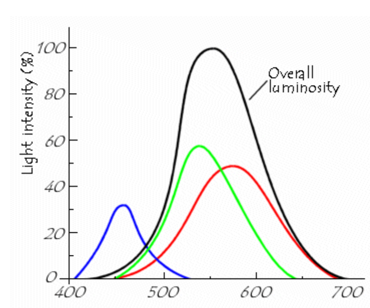

10/7/2019 8

Lossy Image Compression

- Data is lost from the original image to achieve higher compression
- Exploits perceptual redundancy in image data
- Much higher compression ratios can be achieved

  e.g. JPEG can achieve a compression of 100:1

  A color image of this size will require 1920 x 1080 x 3 = ~6.2MB reduced to 62KB

  JPEG – Joint Photographic Encryption Group

JPEG Compression: Loss of Image Quality

Original Image
Size: 219,726 bytes

Highest Quality (Q=100)
Size: 81,447 bytes
Compression: 2.7:1

Medium Quality (Q=25)
Size: 9,407 bytes
Compression: 23:1

Lowest Quality (Q=1)
Size: 1,523 bytes
Compression: 144:1

Image Steganography

Discuss…

Video Compression

- Video Formats

• Standard Definition (SD)	858 x 480	480p
• High Definition (HD)	1280 x 720	720p
• Full-HD	1920 x 1080	1080p
• Ultra HD (UHD)	3840 x 2160	4K
• 8K	7680 x 4320	8K

- Different video streaming services use some, or all of the above formats to stream video.

  i.e. they have to store video in many formats!

Video Compression: Common Formats

- Video Formats
  - Standard Definition (SD) 858 x 480 480p
  - High Definition (HD) 1280 x 720 720p
  - Full-HD 1920 x 1080 1080p
  - Ultra HD (UHD) 3840 x 2160 4K
  - 8K 7680 x 4320 8K
- Audio
  - Most video streams also have an audio component.

Video Compression: Requirements

- A 2-hour Full-HD movie (1080p)

  resolution: 1920 x 1080 pixels/frame = 2,073,600 pixels

  30 frames/second = 2,073,600*30 = 62,208,000 pixels

  3 bytes/pixel (for R, G, B components)

  # bytes of data/second: 1920*1080*30*3 = 186,624,000 bytes/second

  # bytes of data for a 2-hour movie: 186,624,000*60*60*2 = 1,343,692,800,000

Video Compression: Requirements

- # bytes of data/second: 1920*1080*30*3 = 186,624,000 bytes/second
- That is = 1,492,992,000 bits/second

  or ~1.5 GBits/second

- A DVD can store 4.6 Gbytes of data

  Thus a 2-hour movie that takes 1,343,692,800,000 bytes of data for a needs to be compressed ~300:1 to store on a DVD.

Video Compression: Bandwidth

- 1,492,992,000 bits/second (1.5 Gbits/second)

  BMC Ethernet: ~1 Gbits/second

  BMC wifi: 58 Mbits/second

  Verizon 4G LTE: 50 Mbits/second
  (peak, 4-12 MBits/second typical)

Video Compression: Bandwidth

- 1,492,992,000 bits/second (1.5 Gbits/second)

  BMC Ethernet: ~1 Gbits/second 1.5:1

  BMC wifi: 58 Mbits/second 18:1

  Verizon 4G LTE: 50 Mbits/second 20:1
  (peak, 4-12 MBits/second typical) 84:1 .. 250:1

  Ideally 1000:1

  How are such compression ratios achieved?

Video Compression

- Intraframe Compression

  Compress each frame using JPEG. Then stitch frames together.

- Interframe Compression

  Most of the picture is the same in most frames
  Uses predictors, static parts, sample, etc. to compress
  Scene change transitions require coding
  In general, less motion = more compression!

Spiderman

Scene Change

Spiderman

Soundtrack…

Video Compression: Parameters

- Pixels/Frame (SD, HD, $K, etc)
- Frame Rate (24, 30, 60)
- Color Depth
- Length of video
- Amount of motion in video
- Key Frames frequencies
- Constant or variable bitrate streaming
- Buffer Size
- Audio sample rate
- Render quality
- Etc. many other parameters

Containers & Codecs

- **Container**

  Is the file format of a video. This is what the sender sends to the viewer. Besides video it has soundtrack: multiple languages), subtitles (multiple languages), etc. Example formats: OGG, Matroska, AVI, MPEG. Also, Quicktime (.mov), Windows Media Video (.wmv), etc.

- **Codec**

  Are coding and decoding algorithms that actually compress/uncompress the data in a container.
  Example codecs: Xvid, DivX, MPEG-2, H.264, etc.

Video Compression: H.264 Standard

- H.264 *aka* MPEG-4 Part 10, or Advanced Video Coding (MPEG AVC) is the most common video compression standard.
- Capable of delivering HD video at 1.5 MBits/second!
  i.e. streaming can be done over cellular networks (4-12 Mbits/second).
- Most CPUs (e.g. Intel Core i3/i5/i7) have an on-chip hardware full HD H.264 encoder.

Video Compression: H.264 Standard

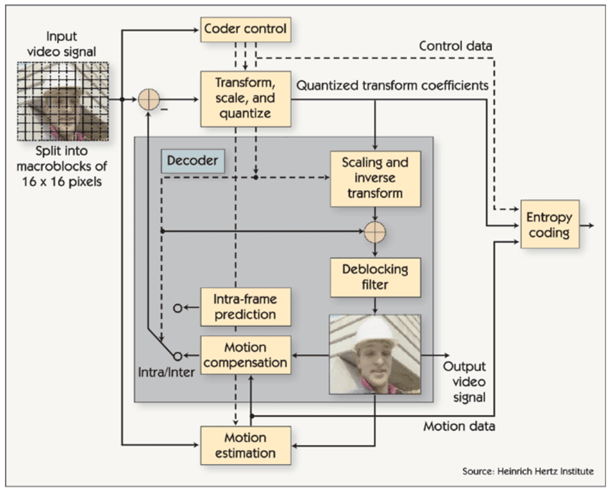

Frontiers of Compression

Frontiers of Compression

- Video streaming makes up 65% of internet traffic
- Netflix and YouTube make up the bulk of it
- At peak download times, Netflix streaming makes up 35% of traffic in USA
- Netflix has 80+ million user accounts
- Over 1 billion hours of video are downloaded from Netflix/week
- Even a tiny bit extra achievement in compression is highly desirable!

Mobile Phone Issues

- Monthly provider data plan limits come into play when streaming movies/video to mobile phone/tablets.
- Mobile phone resolution may not be HD.
- In general, the highest compress factor that can be achieved without sacrificing ***noticeable*** picture quality is the driving factor.
- Unfortunately, lossy compression comes with a loss in picture quality.

Video Compression

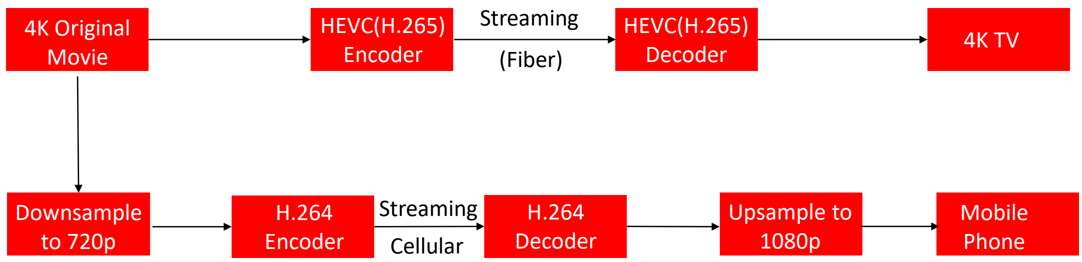

Video Compression

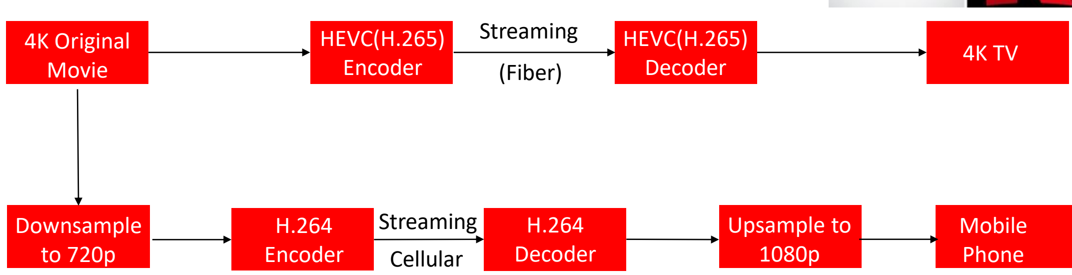

Different codecs are used for different devices, network conditions…
Netflix has to offer different quality levels of the same movie (bitrate ladder)

Goal is to offer best quality for any device and any network condition.

Bitrate Ladder (Netflix prior to 2015)

- Each video is encoded at different bitrates and resolution.

Bitrate [Kbps]	Resolution
1050	640 x 480
1750	720 x 480
2350	1280 x 720
3000	1280 x 720
4300	1920 x 1280
5800	1920 x 1080

VMAF (2015)

- Video Multimethod Assessment Fusion

  Measures human level perceptual quality of video

  Scale from 0 (lowest) to 100 (highest quality)
  Average VMAF is ~70

  Codecs were optimized to use VMAF rating for each movie title
  (Per Title Encoding)

Per Shot Encoding (2018)

- Shots – scenes in film with different video quality
- Algorithms to find best encoding for each shot.

Jessica Jones
Netflix Original Series

Bitrate *vs* Quality Improvements

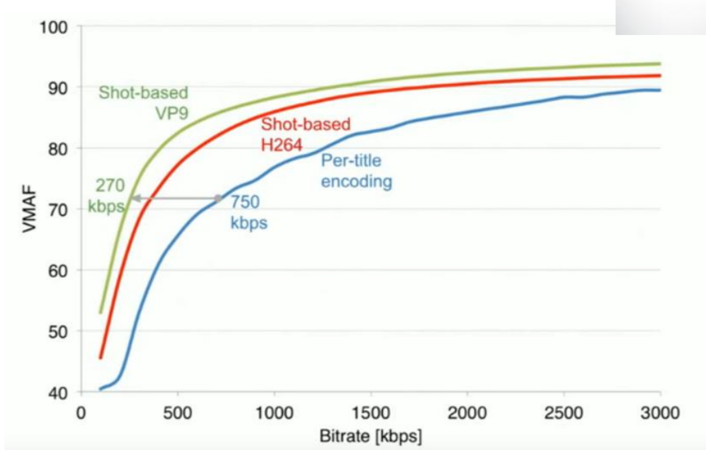

Gains…

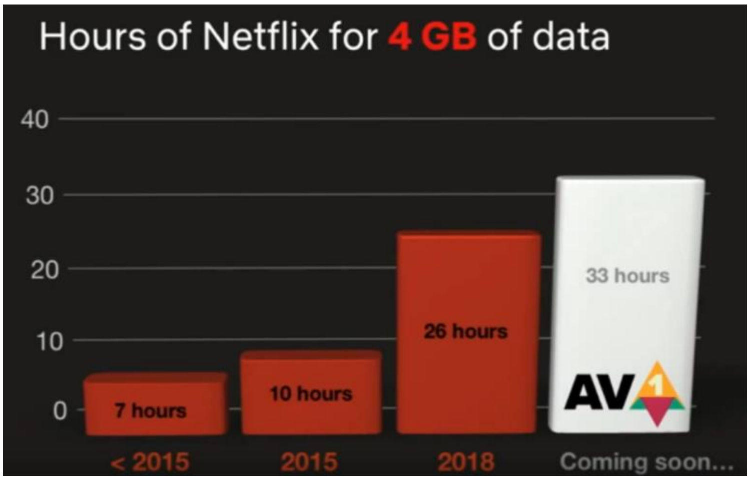

When all else fails…engineer!

- ISPs unhappy with Netflix and streaming services about their use of so much bandwidth
- Netflix offers ISPs free ultra-high efficient storage (100-200TB) and distribution units (OpenConnect Appliances) that store nearly 80% of their streaming library and can be installed in network offices near consumers. This is their Content Delivery Network (CDN).

The Human Experience

- Compression is a technical solution
- Often overuse of compression leads to issues
- Here are some case studies…

Game of Thrones: Battle of Winterfell

"This week's *Game of Thrones* was ostensibly about the Battle of Winterfell, the final confrontation between the armies of the evil Night King and the forces of humanity.

But for many viewers, the episode was much more than an epic battle: It was also a whirlwind tour of the limits of video compression algorithms and home video display technology.

Which is to say that a lot of people couldn't see *anything*."

-: Matthew Dessem, Slate, April 29, 2019.

Game of Thrones: Battle of Winterfell

- Digital film making enables shots in very low light conditions.
- But, switching from CRT TVS to HDTVs renders low light
  detail worse. Especially on cheaper LCD/LED TVs.
- Plus, video compression codecs exacerbate the problem.
  Compression is not good for low light scenes.
- Essentially, advances in technology are the problem!

Compression: Music

- In music production, there is always a testing stage where a mix is played on a car's speakers.
- That is, who will listen to the music on which device is not controllable.
- Goal is to test to ensure that it will sound tolerable/acceptable in all conditions. But this is "lowering the ceiling to raise the floor".
- Same mix, designed for car's speakers may not sound that good on a high-end audio system. And *vice versa*.

Neil Young

Neil Young

He hates Spotify.
He hates Facebook.
He hates Apple.
He hates Steve Jobs.
He hates what digital technology is doing to music.
…
When you hear real music, you get lost in it…because it sounds like God. Spotify doesn’t sound like God. It sounds like a rotating electric fan that someone bought at a hardware store.

-: David Samuels, NY Times, 2019.

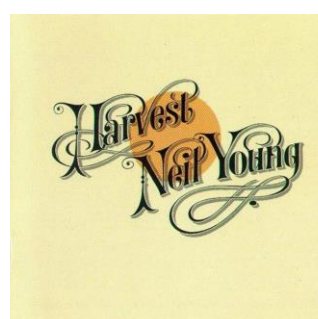

Neil Young

Silicon Valley's emphasis on compression and speed comes at the expense of the notes as they were actually played and is doing something bad to music, which is supposed to make us feel good. It is doing something bad to our brains.

-: David Samuels, NY Times, 2019.

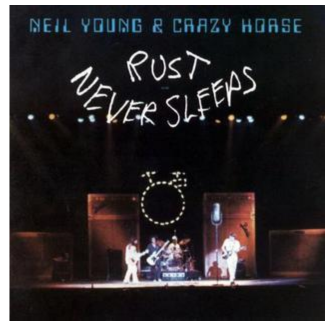

Neil Young

When it comes to listening and enjoying music the internet was as if a meteor had wiped out the existing planet of sound.

The compressed hollow sound of free streaming music was a big step down from the CD. Huge step down from vinyl.

…eliminated levels of sonic detail and shading by squeezing down the amount of information contained in the package in which the music was delivered. Essentially you are left with 5% of the original music for your listening enjoyment.

Neil Young

Listening to music via the current dominant streaming formats is like walking into the Metropolitan Museum of Art or the Musee d'Orsay one morning and finding that all of the great canvases in those museums were gone and replaced by pixelated thumbnails.

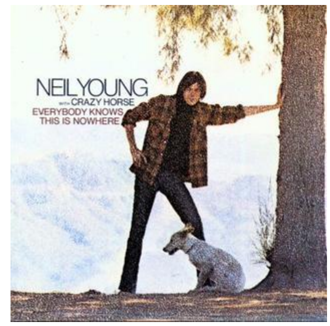

10/7/2019 45

Neil Young

Even engineers in Silicon Valley can hear the difference in the stuff they are selling.

Steve Jobs. He loved music. He listened to vinyl in his living room because he could hear real music.

As, Tim Cook, head of Apple, recently told a reporter, "We worry that that the humanity is being drained out of music."

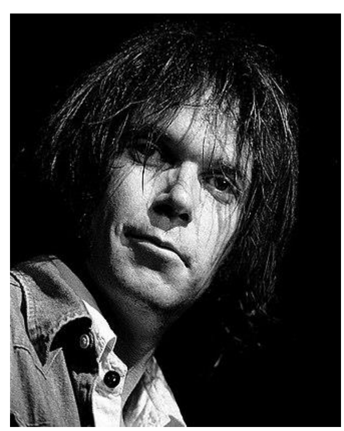

10/7/2019 46

Frontiers of Compression

- **Pattern Matching**

  Given a text pattern P, and a text T
  Find all occurrences of P in T

  E.g.

  $ egrep "Sawyer" TomSawyer.txt

```
Tom Sawyer
of mine. Huck Finn is drawn from life; Tom Sawyer also, but not from an
Sawyer."
when hope was dead, Tom Sawyer came forward with
"Tom Sawyer—Sir,"
Tom Sawyer went home quite cheerful, thinking to himself that there was
MONDAY morning found Tom Sawyer miserable. Monday morning always
…
```


Frontiers of Compression

- **Compressed Pattern Matching**

  Given a text pattern P, and a **compressed** text T
  Find all occurrences of P in T, **without uncompressing**.

  E.g.

  $ egrep "Sawyer" TomSawyer.txt.gz

- Can we do this?

Frontiers of Compression

- **Compressed Pattern Matching**

  Given a text pattern P, and a **compressed** text T
  Find all occurrences of P in T, **without uncompressing**.

  E.g.

  $ egrep “Sawyer” TomSawyer.txt.gz

- Can we do this?
- Why would this be important when one can always uncompress the file and then search?

Frontiers of Compression

- **Compressed Pattern Matching**

  Given a text pattern P, and a **compressed** text T
  Find all occurrences of P in T, **without uncompressing**.

  E.g.

  $ egrep “Sawyer” TomSawyer.txt.gz

- Can we do this?
- Why would this be important when one can always uncompress the file and then search?
- What about things other than text: images, music, etc.?

Frontiers of Compression: Genomic Data

- DNA and RNA sequencing for disease screening is on the rise.
- DNA sequencing of entire populations can give a more complete picture of society-wide health. Recall, Floridi's eHealth reference from Week#1.

  E.g. UK's Biobank project (ukbiobank.ac.uk): sequence the genomes of 500,000 volunteers and track them for decades.
- Sequencing genomes of organisms in the air, soil, water, and even inside our organs can help track epidemics, food toxins, etc.

Frontiers of Compression: Genomic Data

- DNA sequencers that process the human genome generate 10s to 100s of gigabytes of data (e. g. FASTQ files).
- By 2025, it is estimated that 1 billion people would have sequenced their genome.
- This data could easily occupy several exabytes (10^{18}bytes) of storage. We have another data deluge at our hands!
- Current compression algorithms (like gzip) can shrink the size of a genome dataset by a factor of 20.

Growth of Genome Sequencing

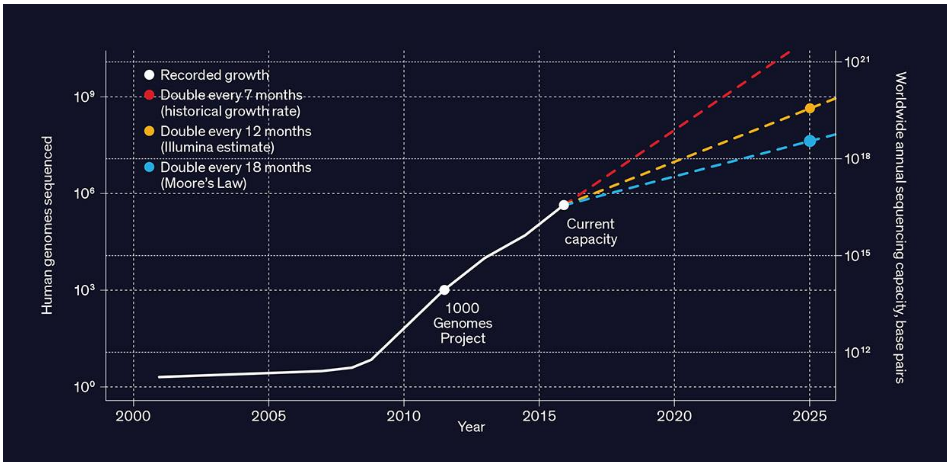

Compressing Genomic Data

- Just like specialized compression techniques for images, video, and audio, there is a need to develop more efficient, faster, compression techniques suited for genomic data.
- Genome data compression has to be lossless. However, some quality data from sequencers can be redundant and so researchers are looking at developing lossy compression schemes.
- Genomic data compression algorithms, just video compression, have to be standardized. MPEG-G is one such specification.

Science Meets Arts Meets Science

Science Meets Art Meets Science…

- The creators of the HBO show **Silicon Valley** tapped Tsachy Weissman (Stanford University) for a technology to feature.
- A *universal compression algorithm:* **a form of powerful and efficient lossless compression that can work on any type of data and is searchable**.
- That is, is it possible to search for features in a compressed DNA without uncompressing every genome in a database?
- A metric, **Weissman Score**, was proposed to compare compression algorithms' performance.

Weissman Score

$$W = \alpha \frac{r}{\overline{r}} \frac{\log \overline{T}}{\log T}$$

NOTE: r and T refer to the compression ratio and time-to-compress for the target algorithm, $\overline{r}$ and $\overline{T}$ refer to the same quantities for a standard universal compressor (e.g. gzip or FLAC), and α is a scaling constant. By normalizing by the performance of a standard compressor, we take away variation in compressive performance between types of data.

References

- JPEG, Wikipedia https://en.wikipedia.org/wiki/JPEG#JPEG_codec_example (10/1/2019)
- Anne Aaron, Netflix. *Tidying up (bits on the internet) with AnneAaron*, Stanford Compression Forum, 2019. Video of talk is available at: https://www.youtube.com/watch?v=xELiq0HtuzE&feature=youtu.be&list=PLv_7iO_xlL0JOUKrZSSu0k19gl5Jh9KDd (10/4/2019)
- Matthew Dessem, *Why You Couldn't See a Damn Thing on This Week's Game of Thrones*, Slate, April 29, 2019. https://slate.com/culture/2019/04/game-of-thrones-long-night-dark-cinematography-twitter-complaints.html
- David Samuels. *Neil Young's Lonely Quest to Save Music*. New York Times, August 20, 2019.
- Dmitri Pavlichin & Tsachy Weissman. *The Desperate Quest for Genomic Compression Algorithms*. IEEE Spectrum, 22 August, 2018. Available at: https://spectrum.ieee.org/computing/software/the-desperate-quest-for-genomic-compression-algorithms
- *What Happens When you Press Play*. High Scalability. 2017. http://highscalability.com/blog/2017/12/11/netflix-what-happens-when-you-press-play.html?currentPage=2