Lecture 5: Privacy Engineering Tools and Approaches

CMSC 25910

Winter 2026

The University of Chicago

THE UNIVERSITY OF CHICAGO

Implementing Key Data Privacy Rights

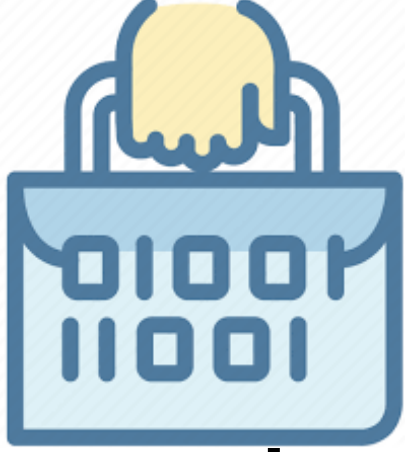

Art. 20 GDPR

Right to data portability

1. The data subject shall have the right to receive the personal data concerning him or her, which he or she has provided to a controller, in a structured, commonly used and machine-readable format and have the right to transmit those data to another controller without hindrance from the controller to which the personal data have been provided, where:

Art. 15 GDPR

Right of access by the data subject

The data subject shall have the right to obtain from the controller confirmation as to whether or not personal data concerning him or her are being processed, and, where that is the case, access to the personal data and the following information:

[1]The controller shall take appropriate measures to provide any information referred to in Articles 13 and 14 and any communication under Articles 15 to 22 and 34 relating to processing to the data subject in a concise, transparent, intelligible and easily accessible form, using clear and plain language, in particular for any information addressed specifically to a child. [2]The information shall be provided in writing, or by other means, including, where appropriate, by electronic means. [3]When requested by the data subject, the information may be provided orally, provided that the identity of the data subject is proven by other means.

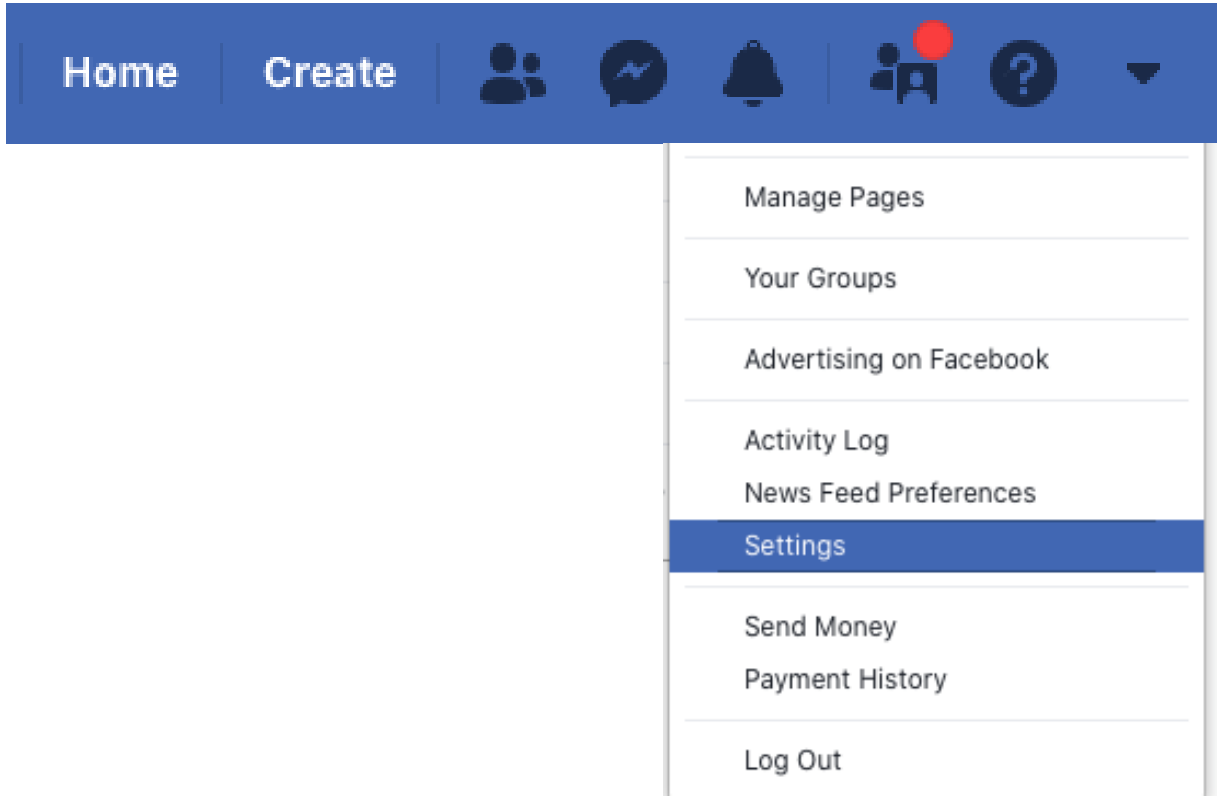
Home
Create
Manage Pages
Your Groups
Advertising on Facebook
Activity Log
News Feed Preferences
Settings
Send Money
Payment History
Log Out

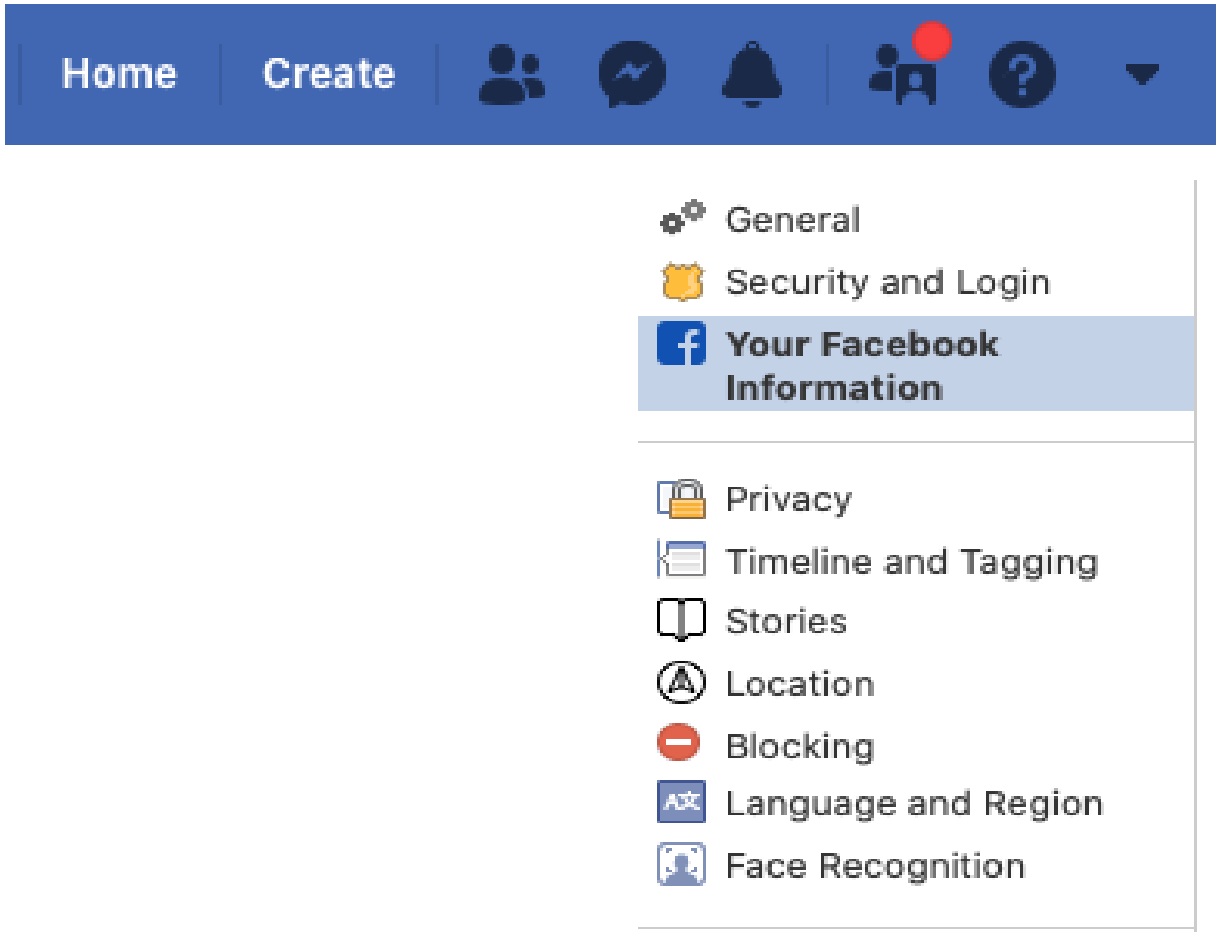
Home
Create
General
Security and Login
Your Facebook Information
Privacy
Timeline and Tagging
Stories
Location
Blocking
Language and Region
Face Recognition

Access Your Information	View your information by category.	View
Download Your Information	Download a copy of your information to keep, or to transfer to another service.	View
Activity Log	View and manage your information and some settings.	View
Off-Facebook Activity	View or clear activity from businesses and organizations you visit off of Facebook.	View
Managing Your Information	Learn more about how you can manage your information.	View
Deactivation and Deletion	Temporarily deactivate or permanently delete your account.	View

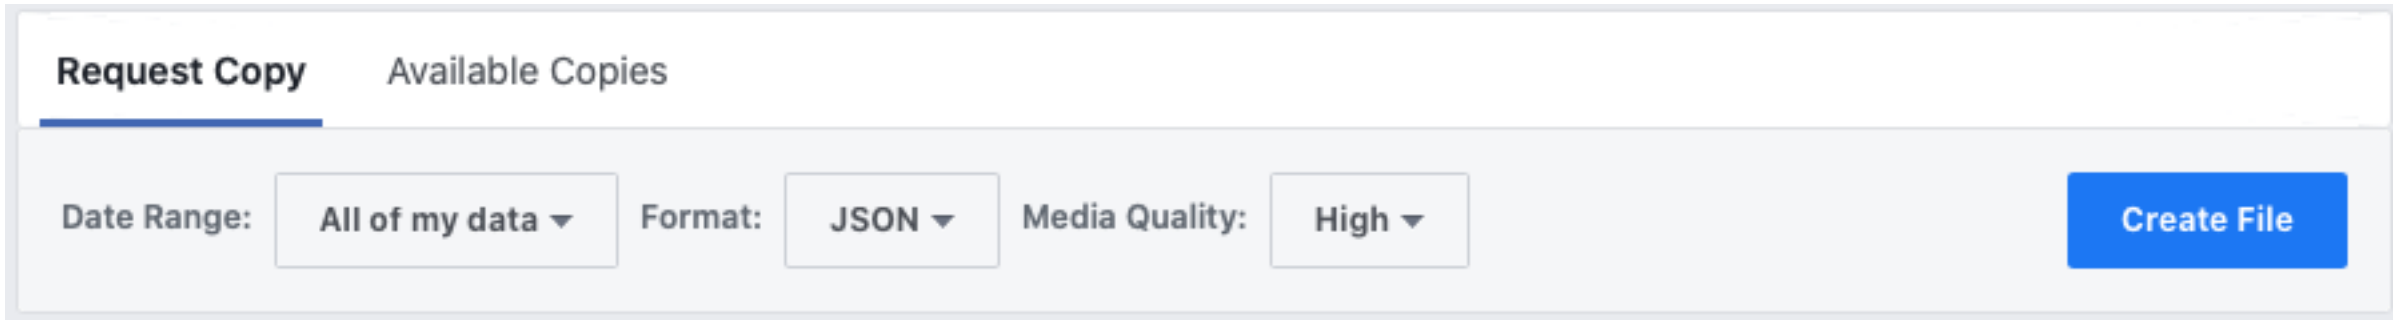
Request Copy
Available Copies
Date Range:
All of my data
Format:
JSON
Media Quality:
High
Create File

A copy of your information is being created.

Your copy may contain more than one file, depending on how much information your request contains. We'll let you know when your copy is complete, so you can download it to your preferred device. You can cancel this process before the file is complete.

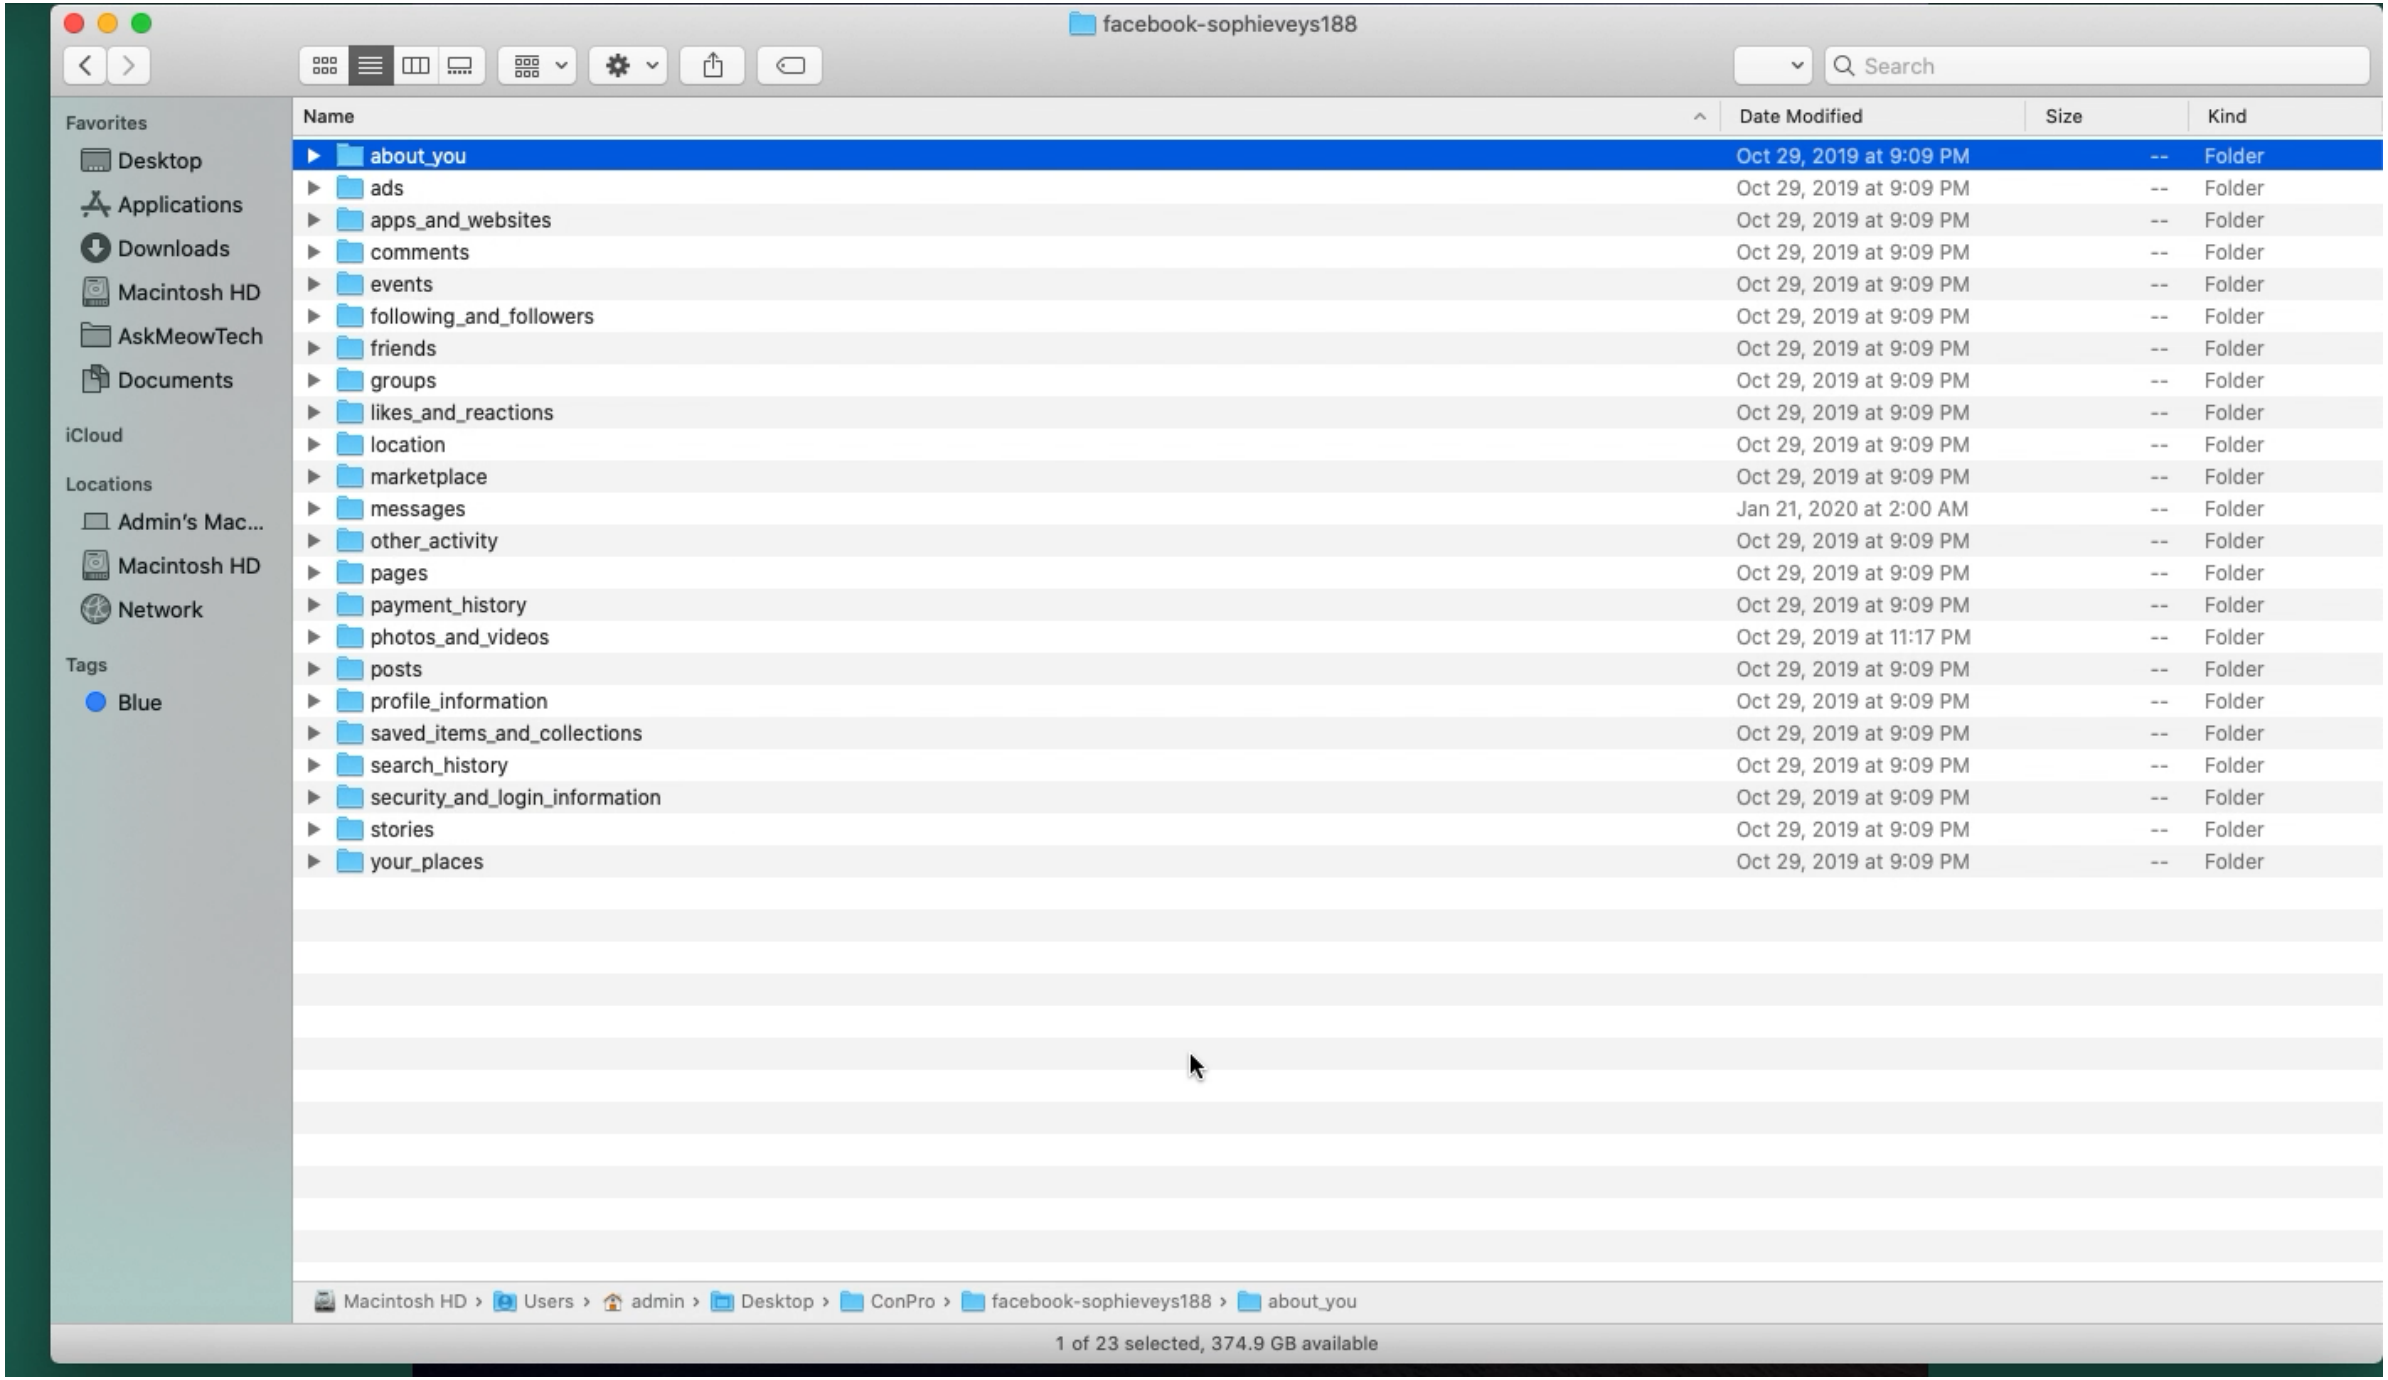
facebook-sophieveys188
Search
Favorites
Desktop
Applications
Downloads
Macintosh HD
AskMeowTech
Documents
iCloud
Locations
Admin's Mac...
Macintosh HD
Network
Tags
Blue
Name
Date Modified
Size
Kind
about_you Oct 29, 2019 at 9:09 PM -- Folder
ads Oct 29, 2019 at 9:09 PM -- Folder
apps_and_websites Oct 29, 2019 at 9:09 PM -- Folder
comments Oct 29, 2019 at 9:09 PM -- Folder
events Oct 29, 2019 at 9:09 PM -- Folder
following_and_followers Oct 29, 2019 at 9:09 PM -- Folder
friends Oct 29, 2019 at 9:09 PM -- Folder
groups Oct 29, 2019 at 9:09 PM -- Folder
likes_and_reactions Oct 29, 2019 at 9:09 PM -- Folder
location Oct 29, 2019 at 9:09 PM -- Folder
marketplace Oct 29, 2019 at 9:09 PM -- Folder
messages Jan 21, 2020 at 2:00 AM -- Folder
other_activity Oct 29, 2019 at 9:09 PM -- Folder
pages Oct 29, 2019 at 9:09 PM -- Folder
payment_history Oct 29, 2019 at 9:09 PM -- Folder
photos_and_videos Oct 29, 2019 at 11:17 PM -- Folder
posts Oct 29, 2019 at 9:09 PM -- Folder
profile_information Oct 29, 2019 at 9:09 PM -- Folder
saved_items_and_collections Oct 29, 2019 at 9:09 PM -- Folder
search_history Oct 29, 2019 at 9:09 PM -- Folder
security_and_login_information Oct 29, 2019 at 9:09 PM -- Folder
stories Oct 29, 2019 at 9:09 PM -- Folder
your_places Oct 29, 2019 at 9:09 PM -- Folder
Macintosh HD › Users › admin › Desktop › ConPro › facebook-sophieveys188 › about_you
1 of 23 selected, 374.9 GB available

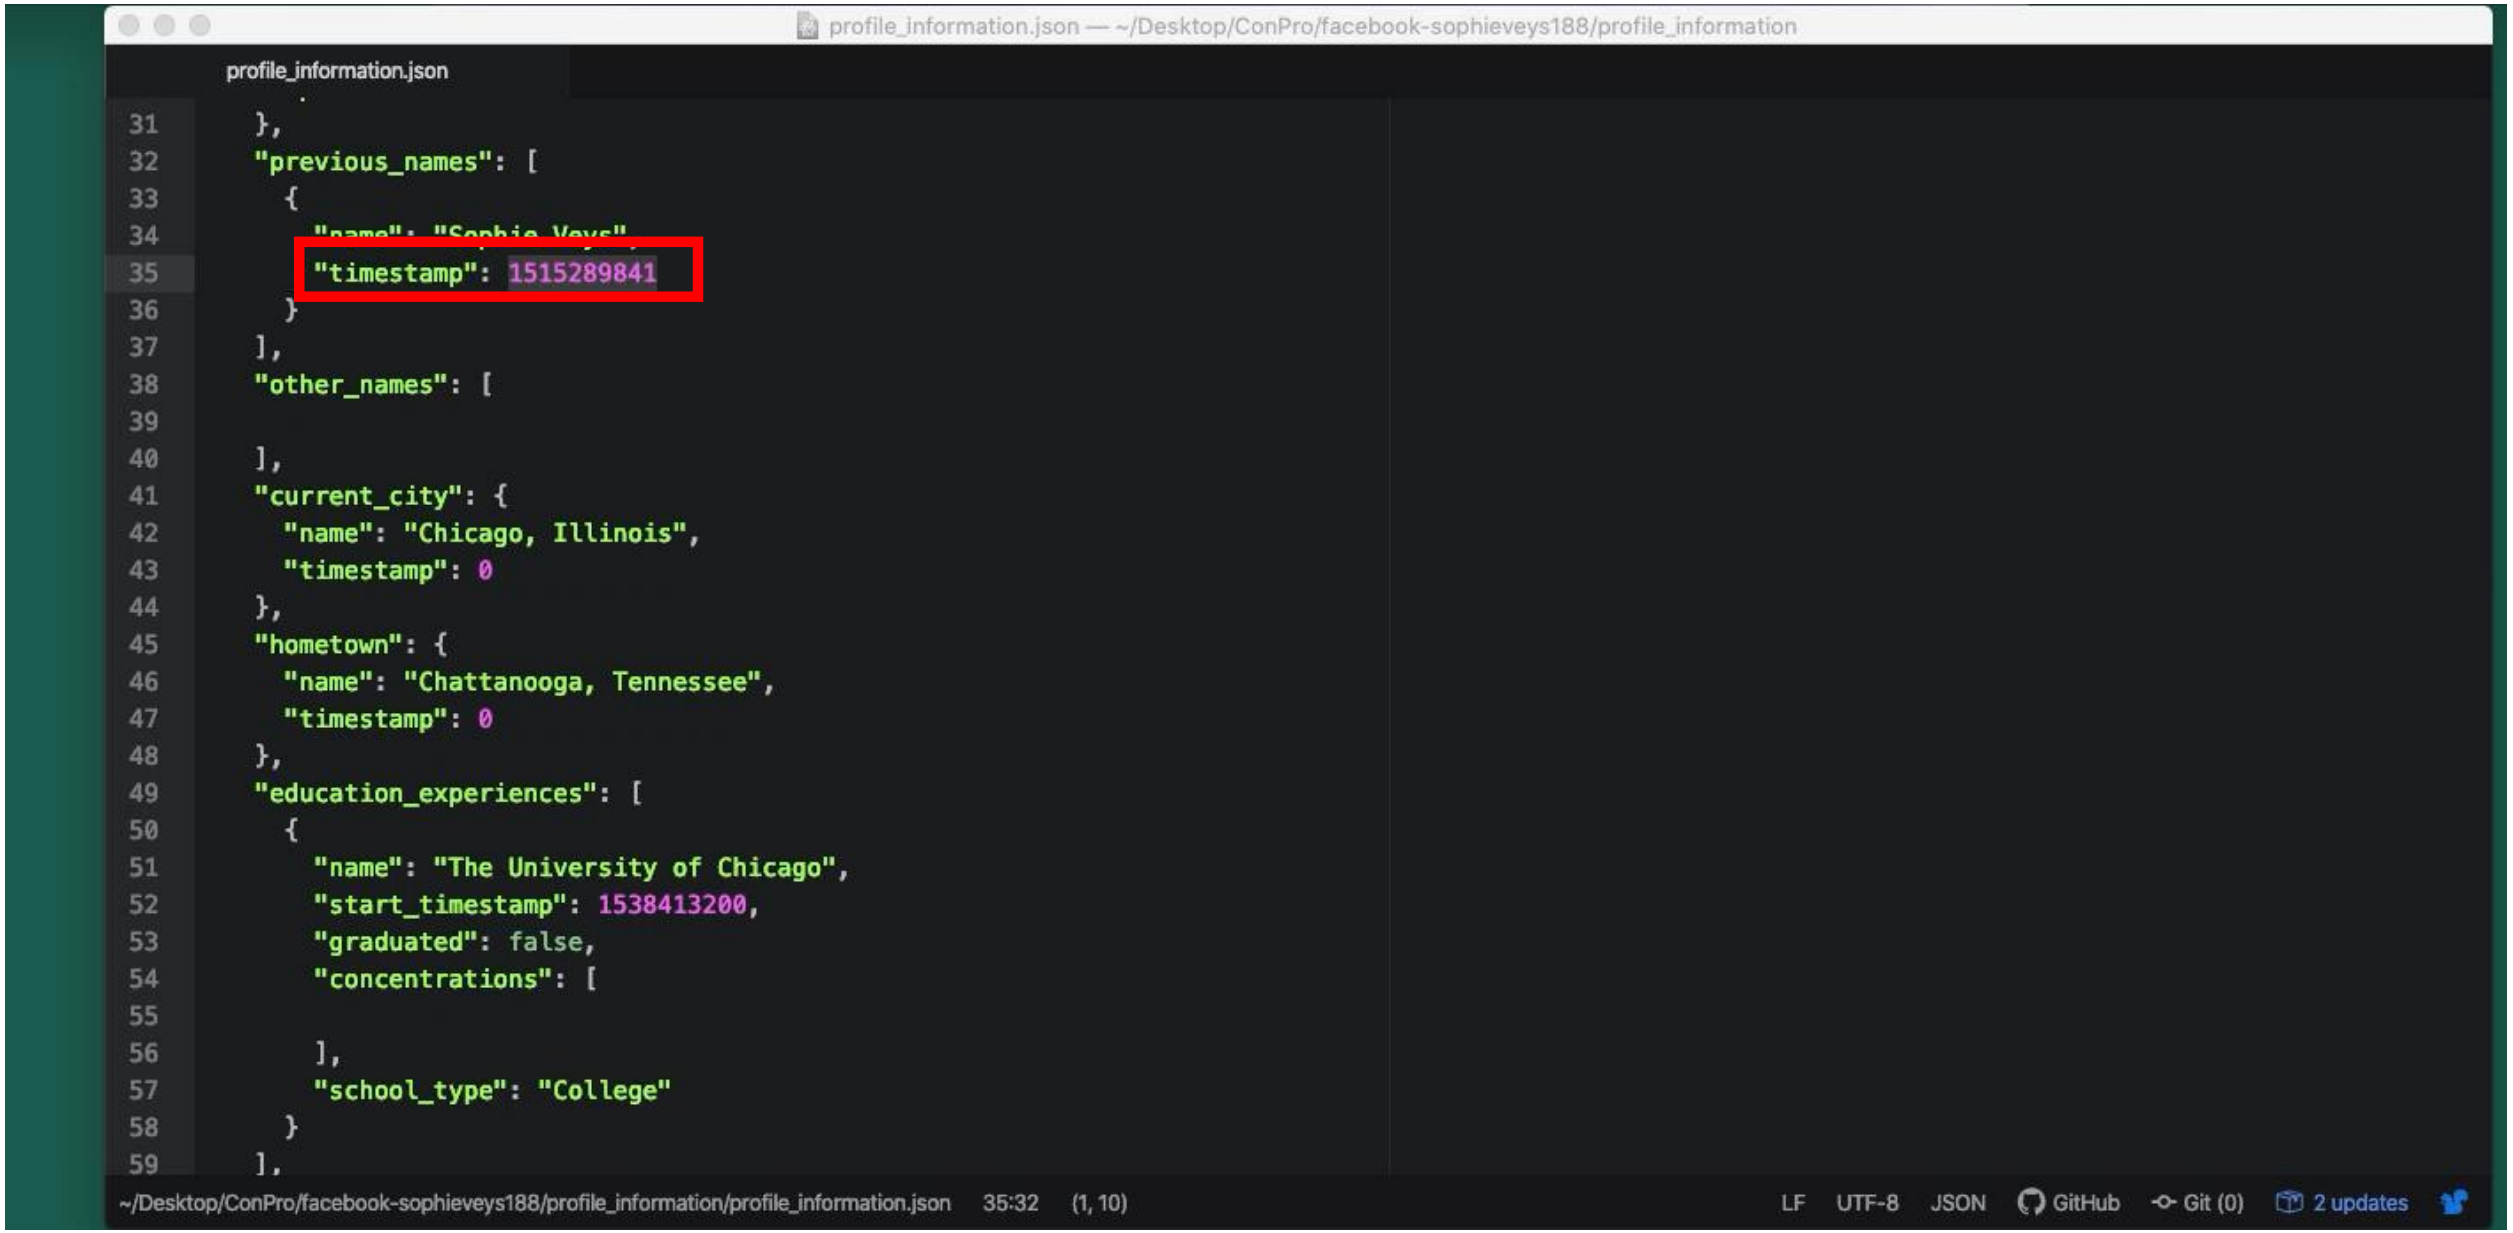
profile_information.json — ~/Desktop/ConPro/facebook-sophieveys188/profile_information
profile_information.json
},
"previous_names": [
{
"name": "Sophie Veys",
"timestamp": 1515289841
}
],
"other_names": [
],
"current_city": {
"name": "Chicago, Illinois",
"timestamp": 0
},
"hometown": {
"name": "Chattanooga, Tennessee",
"timestamp": 0
},
"education_experiences": [
{
"name": "The University of Chicago",
"start_timestamp": 1538413200,
"graduated": false,
"concentrations": [
],
"school_type": "College"
}
],
~/Desktop/ConPro/facebook-sophieveys188/profile_information/profile_information.json 35:32 (1, 10)
LF UTF-8 JSON GitHub Git (0) 2 updates

```
"advertiserInfo" : {
  "advertiserName" : "Foundation Medicine",
  "screenName" : "@FoundationATCG"
},
"matchedTargetingCriteria" : [ {
  "targetingType" : "Keywords",
  "targetingValue" : "#ASCO19"
}, {
  "targetingType" : "Locations",
  "targetingValue" : "United States"
} ],
"impressionTime" : "2019-06-01 08:21:18"
```



```
{
  "sender_name": "Bob",
  "timestamp_ms": 1569679976090,
  "content": "Video call DND session?",
  "type": "Generic"
},
```


Engineering Challenges In Supporting Data Access/Portability

Challenges Implementing Data Access

- Authenticating data subject access requests (DSARs)

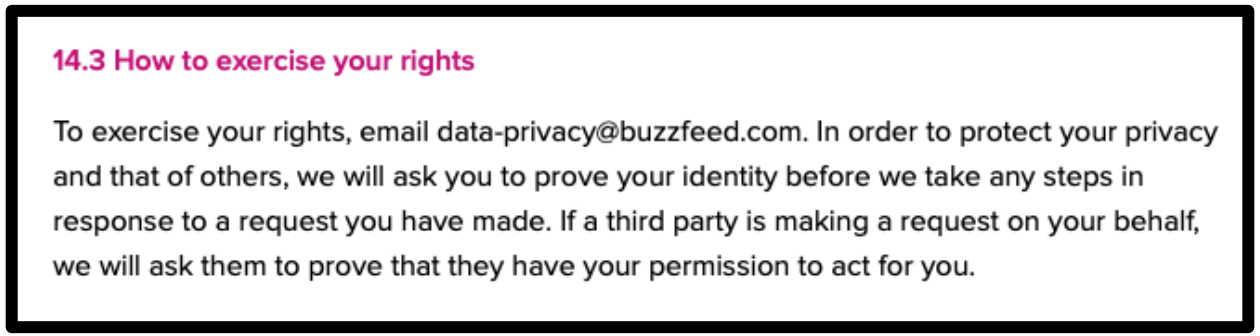

14.3 How to exercise your rights

To exercise your rights, email data-privacy@buzzfeed.com. In order to protect your privacy and that of others, we will ask you to prove your identity before we take any steps in response to a request you have made. If a third party is making a request on your behalf, we will ask them to prove that they have your permission to act for you.

- Deciding what data should be included (e.g., derived data?)
- Finding all data that is included
- What should the format of the response be?
  - What definitions are needed as part of **data dictionaries**
- What should the user experience be?

Right to Erasure

GDPR Article 17

- **Right to erasure (formerly known as the “right to be forgotten”)**
- The data subject shall have the right to obtain from the controller the **erasure of personal data concerning him or her** without undue delay and the controller shall have the obligation to erase personal data without undue delay where one of the following grounds applies:
  - the **personal data are no longer necessary in relation to the purposes for which they were collected** or otherwise processed;
  - the data subject **withdraws consent**…and where there is no other legal ground for the processing;
  - the **data subject objects to the processing…**and there are no overriding legitimate grounds for the processing

CCPA (California Civic Code 1798.105)

- (a) A consumer shall have the **right to request that a business delete any personal information about the consumer which the business has collected from the consumer**…
- (c) A business that receives a verifiable consumer request from a consumer to delete the consumer's personal information pursuant to subdivision (a) of this section shall delete the consumer's personal information from its records and direct any service providers to delete the consumer's personal information from their records.
- (d) A business or a service provider shall not be required to comply with a consumer's request to delete the consumer's personal information if it is necessary for the business or service provider to maintain the consumer's personal information in order to:
  - (1) **Complete the transaction** for which the personal information was collected, fulfill the terms of a written warranty or product recall conducted in accordance with federal law, provide a good or service requested by the consumer, or reasonably anticipated within the context of a business' ongoing business relationship with the consumer, or otherwise perform a contract between the business and the consumer.
  - (2) **Detect security incidents, protect against malicious, deceptive, fraudulent, or illegal activity**; or prosecute those responsible for that activity.
  - (3) **Debug to identify and repair errors** that impair existing intended functionality…
  - (6) Engage in public or peer-reviewed scientific, historical, or statistical research in the public interest that adheres to all other applicable ethics and privacy laws, when the business' deletion of the information is likely to render impossible or seriously impair the achievement of such research, if the consumer has provided informed consent.
  - (7) To enable solely internal uses that are **reasonably aligned with the expectations of the consumer** based on the consumer's relationship with the business

Right to Erasure

- Bound to jurisdictions
- Not applicable to public figures / issues of public interest

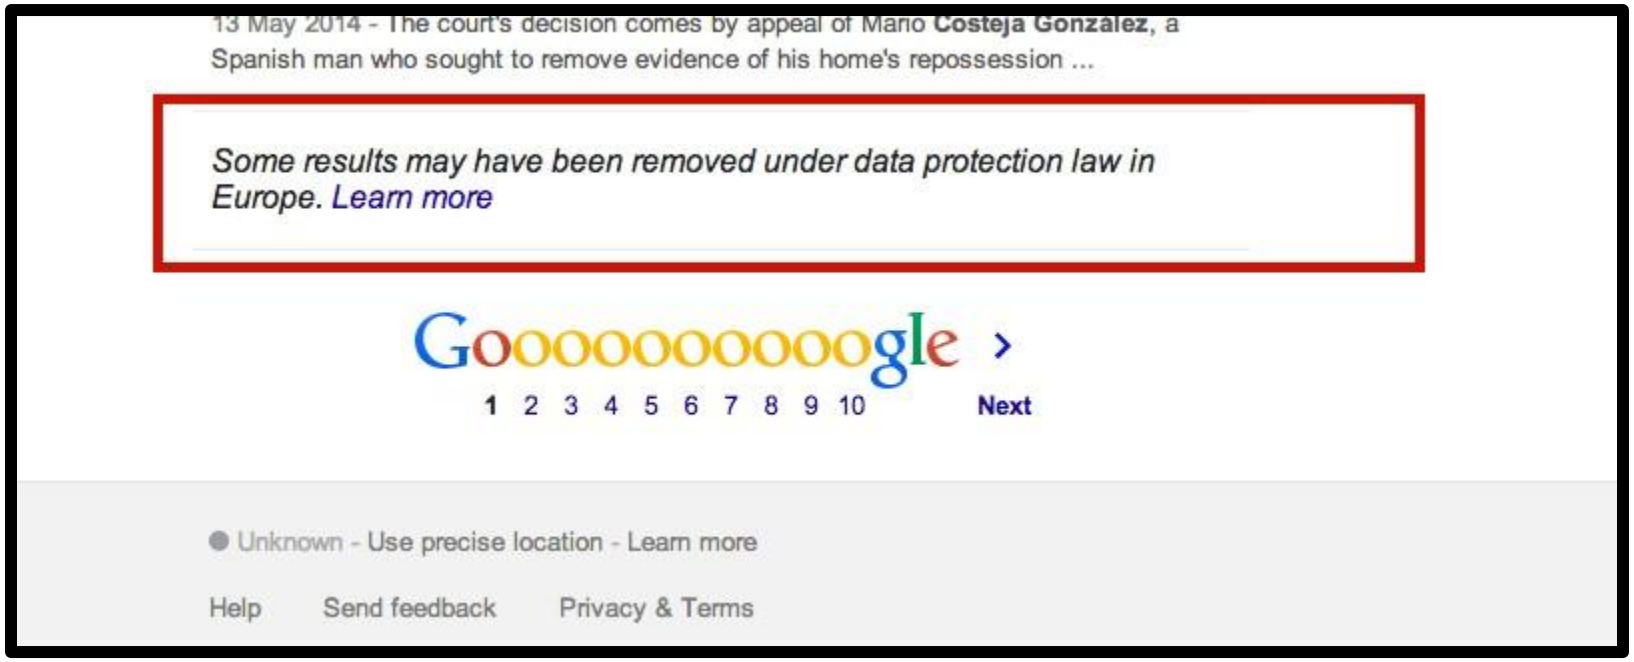

The Difficulty of Data Deletion

For Users: Making Others Delete Data

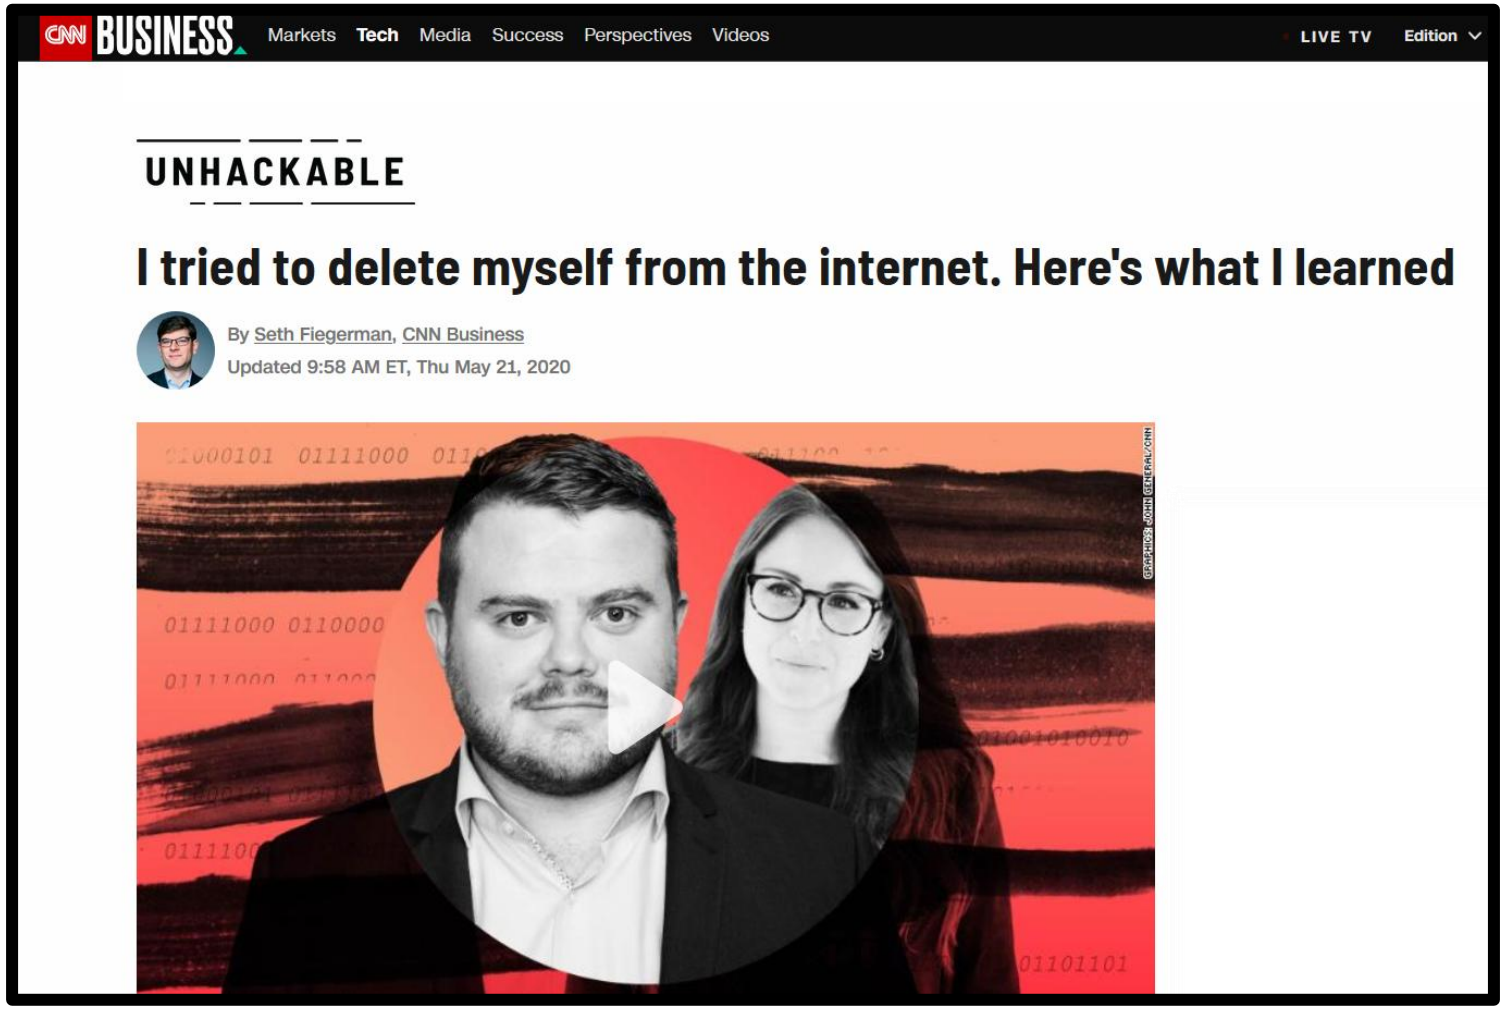

https://www.cnn.com/2020/05/21/tech/deleting-personal-data-online/index.html

For Users: Learning Who Has Your Data

The New York Times

I Got Access to My Secret Consumer Score. Now You Can Get Yours, Too.

Little-known companies are amassing your data — like food orders and Airbnb messages — and selling the analysis to clients. Here's how to get a copy of what they have on you.

https://www.nytimes.com/2019/11/04/business/secret-consumer-score-access.html

Engineering Challenges for Data Deletion

- Backup copies and duplicated data
- Considering log files and other subtle records
- Removing the data that should be erased from aggregated data
  - How does this impact an ML model?
- Legal, contractual, or policy considerations for data *retention*
- Making decisions about “shared” data
- Provenance tracking
  - How do you store metadata about the subject of particular data? How does this vary based on the data structure / data type?

The Nuances of Deletion

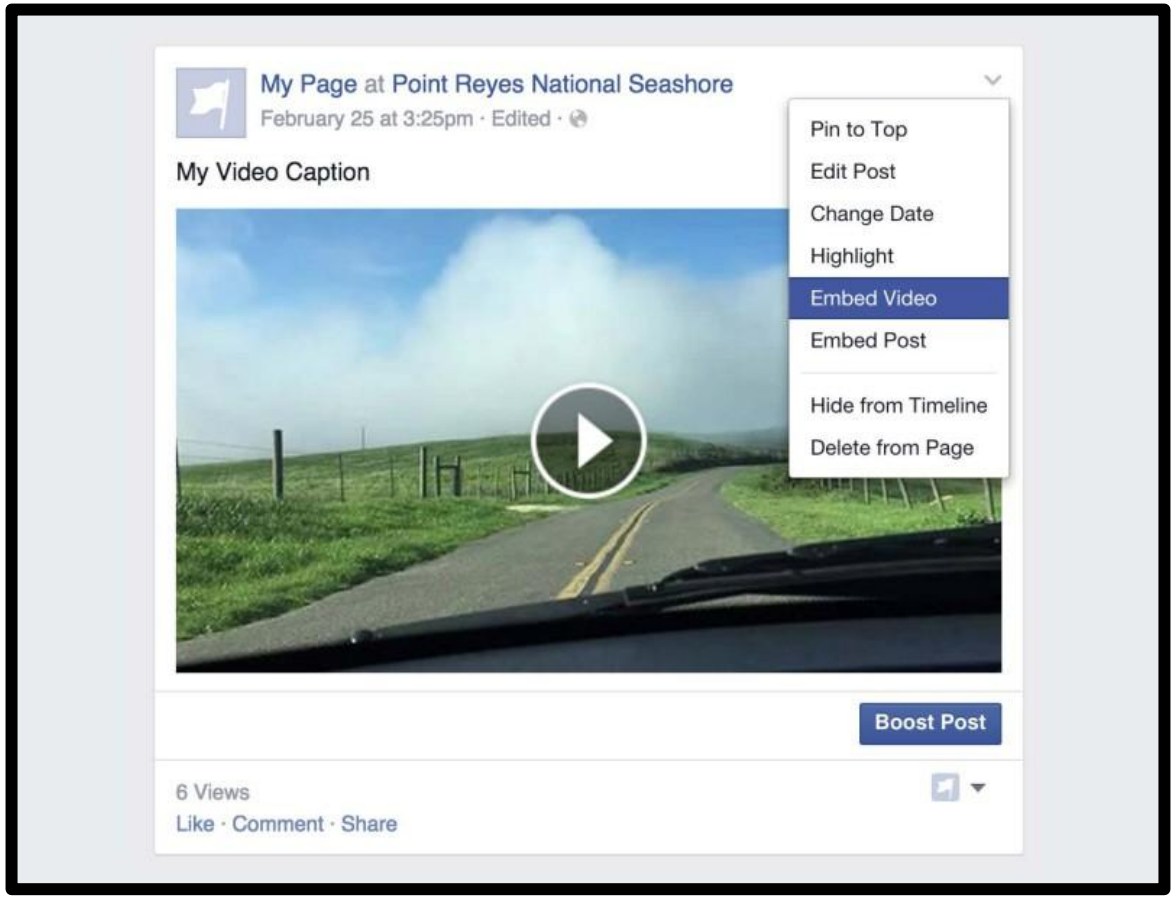

https://marketingland.com/facebook-launches-new-video-embeds-comment-syncing-from-site-to-page-122868

Derived Data Is Challenging

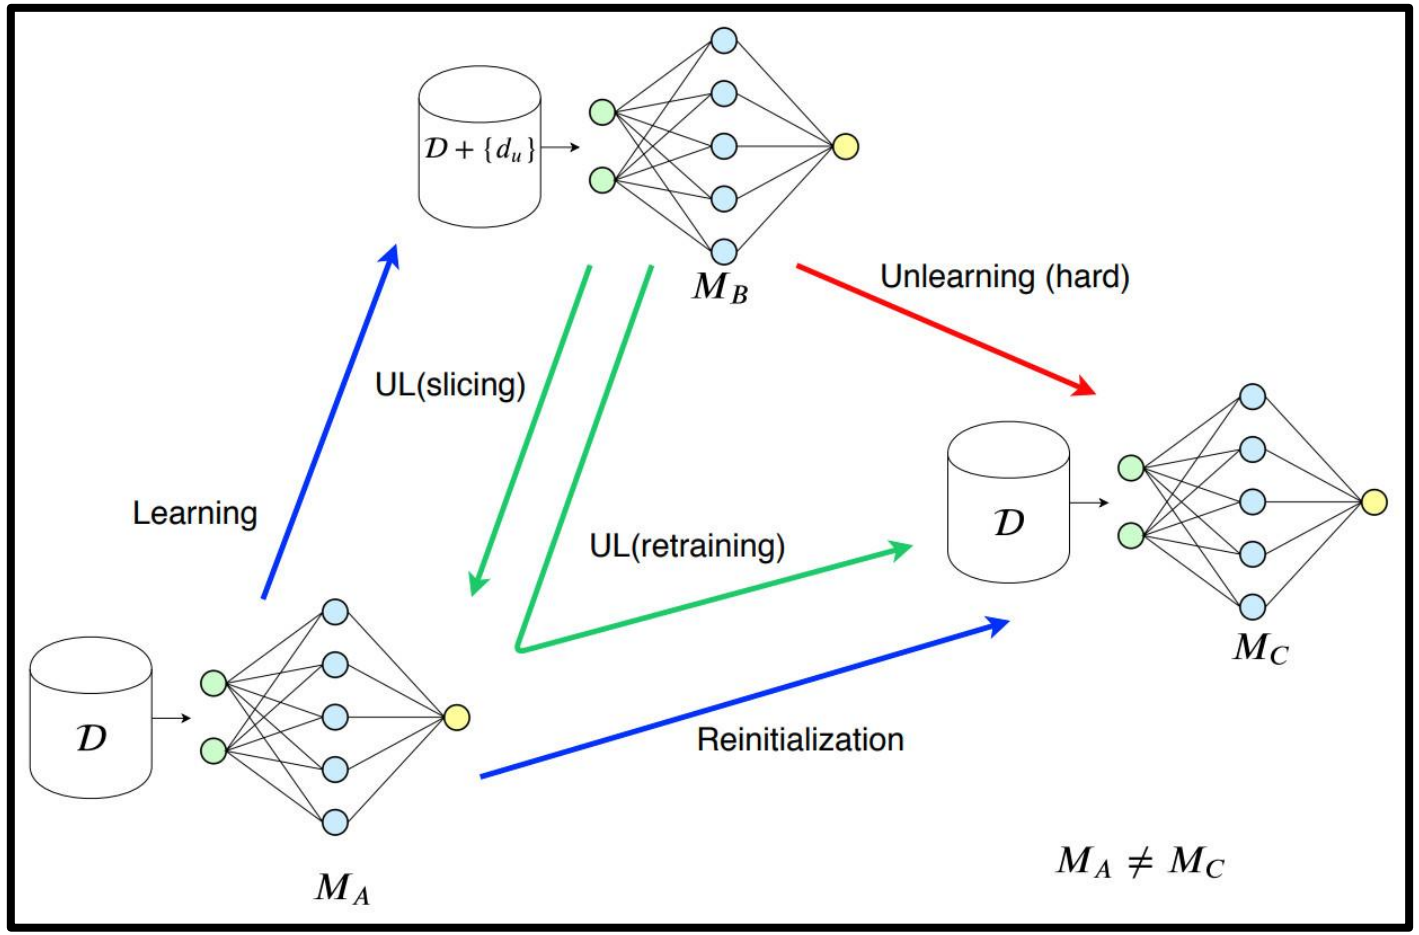

Images taken from https://arxiv.org/abs/1912.03817
The work is summarized on http://www.cleverhans.io/2020/07/20/unlearning.html

Detecting Child Sexual Abuse Material (CSAM) in a Privacy-Preserving Way

Apple’s Proposals For CSAM Detection

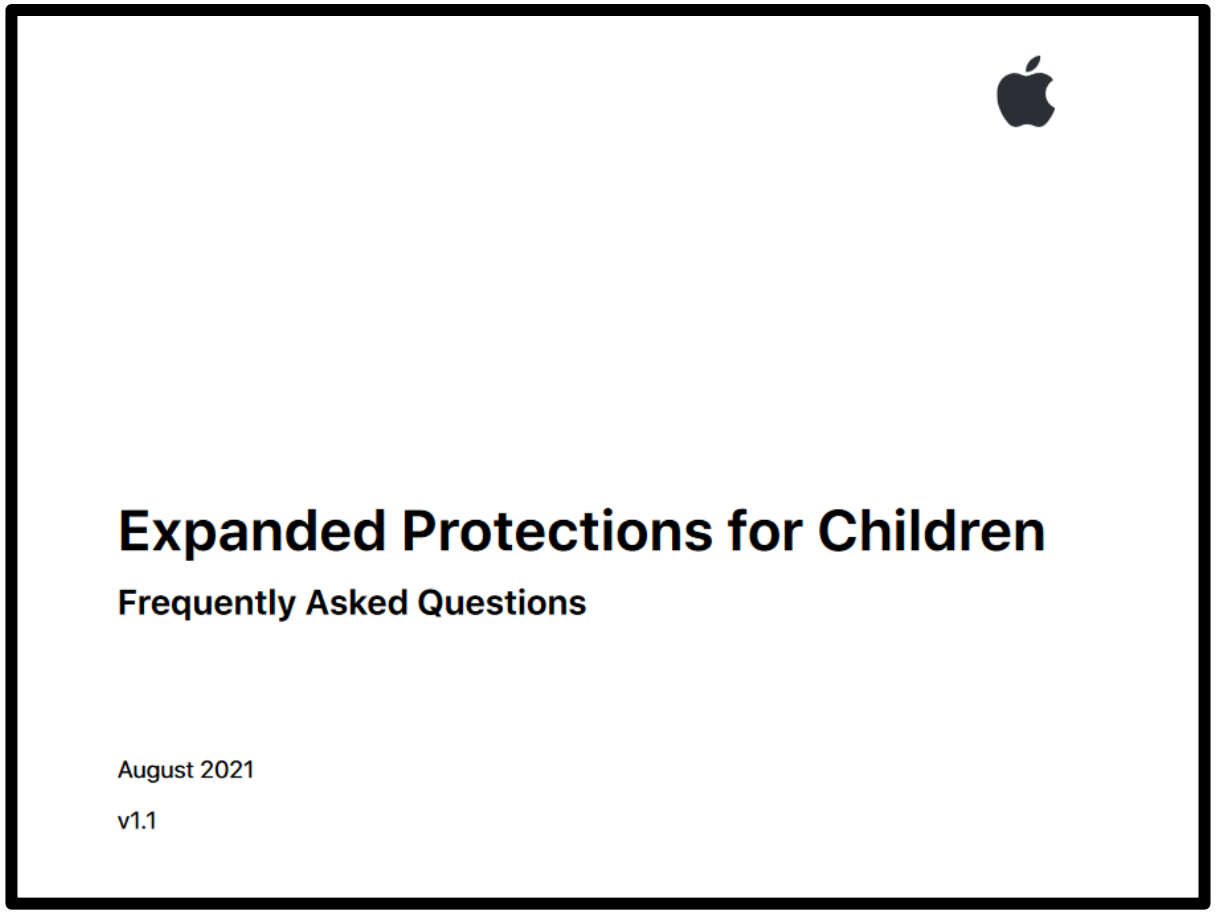

https://www.apple.com/child-safety/pdf/Expanded_Protections_for_Children_Frequently_Asked_Questions.pdf

Class Discussion: Design Requirements

- High-level goal: Be able to detect known CSAM images/videos
- What are the key engineering design requirements?
- What are the key engineering challenges?
- What are the key non-engineering challenges?

Apple’s Proposals For CSAM Detection

- “Apple does not learn anything about images that do not match the known CSAM database.”
- “Apple can’t access metadata or visual derivatives for matched CSAM images until a threshold of matches is exceeded for an iCloud Photos account.”
- “The risk of the system incorrectly flagging an account is extremely low. In addition, Apple manually reviews all reports made to NCMEC to ensure reporting accuracy.”
- “Users can’t access or view the database of known CSAM images.”
- “Users can’t identify which images were flagged as CSAM by the system.”

All points taken verbatim from https://www.apple.com/child-safety/pdf/CSAM_Detection_Technical_Summary.pdf

Apple’s Proposals For CSAM Detection

- “Instead of scanning images in the cloud, the system performs on-device matching using a database of known CSAM image hashes provided by NCMEC and other child-safety organizations. Apple further transforms this database into an unreadable set of hashes, which is securely stored on users’ devices.”

All points taken verbatim from https://www.apple.com/child-safety/pdf/CSAM_Detection_Technical_Summary.pdf

Detour: Hash Function

- **Function mapping arbitrarily sized data to a fixed output**
  - **Deterministic**: always maps a given input to the same output
- (Usually) desirable properties of hash functions:
  - **Efficiently** computable (not always desirable; see password storage)
  - Aim to minimize collisions (different inputs → same output), though collisions are unavoidable by definition (fixed space of outputs)
  - Aim to map inputs relatively **uniformly** to the space of possible outputs
- **Non-cryptographic hash functions** are often used for hashmaps (e.g., Python dictionaries) and checksums (non-adversarial cases)
- **Cryptographic hash functions** are one-way functions (very hard to invert); small differences in input → large differences in output

Non-Cryptographic Hash Functions

Non-cryptographic hash functions [edit]

Main article: Non-cryptographic hash function

Name	Length	Type [hide]
Pearson hashing	8 bits (or more)	XOR/table
Paul Hsieh's SuperFastHash[1]	32 bits	
Buzhash	variable	XOR/table
Fowler–Noll–Vo hash function (FNV hash)	32, 64, 128, 256, 512, or 1024 bits	xor/product or product/XOR
Jenkins hash function	32 or 64 bits	XOR/addition
Bernstein's hash *djb2*[2]	32 or 64 bits	shift/add or mult/add or shift/add/xor or mult/xor
PJW hash / ELF hash	32 or 64 bits	add,shift,xor
MurmurHash	32, 64, or 128 bits	product/rotation
Fast-Hash[3]	32 or 64 bits	xorshift operations
SpookyHash	32, 64, or 128 bits	see Jenkins hash function
CityHash[4]	32, 64, 128, or 256 bits	
FarmHash[5]	32, 64 or 128 bits	
MetroHash[6]	64 or 128 bits	
numeric hash (nhash)[7]	variable	division/modulo
xxHash[8]	32, 64 or 128 bits	product/rotation

Table from https://en.wikipedia.org/wiki/List_of_hash_functions#Non-cryptographic_hash_functions

Cryptographic Hash Functions

- Less focus on efficiency; focus on resisting **preimage attacks** (given an output, find an input that hashes to that output)

Table from https://en.wikipedia.org/wiki/List_of_hash_functions

Unkeyed cryptographic hash functions [edit]

Main article: Cryptographic hash function
See also: Comparison of cryptographic hash functions

Name	Length	Type [hide]
BLAKE-256	256 bits	HAIFA structure[17]
BLAKE-512	512 bits	HAIFA structure[17]
BLAKE2s	up to 256 bits	HAIFA structure[17]
BLAKE2b	up to 512 bits	HAIFA structure[17]
BLAKE2X	arbitrary	HAIFA structure,[17] extendable-output functions (XOFs) design[18]
BLAKE3	arbitrary	Merkle tree
ECOH	224 to 512 bits	hash
FSB	160 to 512 bits	hash
GOST	256 bits	hash
Grøstl	up to 512 bits	hash
HAS-160	160 bits	hash
HAVAL	128 to 256 bits	hash
JH	224 to 512 bits	hash
LSH[19]	256 to 512 bits	wide-pipe Merkle–Damgård construction
MD2	128 bits	hash
MD4	128 bits	hash
MD5	128 bits	Merkle–Damgård construction
MD6	up to 512 bits	Merkle tree NLFSR (it is also a keyed hash function)
RadioGatún	arbitrary	ideal mangling function
RIPEMD	128 bits	hash
RIPEMD-128	128 bits	hash
RIPEMD-160	160 bits	hash
RIPEMD-256	256 bits	hash
RIPEMD-320	320 bits	hash
SHA-1	160 bits	Merkle–Damgård construction
SHA-224	224 bits	Merkle–Damgård construction
SHA-256	256 bits	Merkle–Damgård construction
SHA-384	384 bits	Merkle–Damgård construction
SHA-512	512 bits	Merkle–Damgård construction
SHA-3 (subset of Keccak)	arbitrary	sponge function

Caution: Hashing For Privacy

- Simply hashing PII seems like it would provide privacy…
  - … but attackers can **enumerate and hash possible inputs of interest!**
- Salt: random value added to input (e.g., prepended, appended) before hashing; the same input with a different salt maps to a different output
- Sometimes, you may want to salt inputs and then discard the salt
  - Protect privacy while recording similarity
  - The salt needs to be big enough to resist brute-forcing

Other Types of Hashing

- **Locality-sensitive hashing**: fuzzy hashing that tries to tries to map similar items to the same bucket
  - Very different than cryptographic hashing despite the name; here, we instead want similar inputs to map to **similar** outputs
  - This often requires you to featurize an image
  - Consider the color distributions, objections recognized, sub-images, and other ways you could partition an image (or other item)
- **Minhash** (min-wise independent permutations locality sensitive hashing scheme) is one type

See https://en.wikipedia.org/wiki/MinHash

Other Types of Hashing

- **Perceptual hashing**: a type of locality-sensitive hashing that tries to model how a human would perceive a media object (e.g., picture, audio, video) as its measure of similarity
- Apple's CSAM proposal used Neural Hash, a specific perceptual hash function (seemingly since abandoned)

See https://www.apple.com/child-safety/pdf/CSAM_Detection_Technical_Summary.pdf

Additional Techniques Used

- **Threshold Secret Sharing:** Cryptographic technique that requires a specified number of shares to reconstruct a secret
- **Private Set Intersection (PSI):** Cryptographic technique that enables two parties to find the intersection of their sets without revealing elements that are not in common

See https://www.apple.com/child-safety/pdf/CSAM_Detection_Technical_Summary.pdf

Detour: Secret Sharing

- (Adi) Shamir's Secret Sharing scheme is based on polynomial interpolation over finite fields
- Insight: k points uniquely determine a polynomial of degree k-1
  - 2 points uniquely define a line; 3 points uniquely define a polynomial of degree 2 ($ax^2 + bx + c$)
- Approach: pick a polynomial of appropriate degree such that the y-intercept is the "secret" and give everyone a point on this polynomial as their "share"
  - Challenge: Prevent Sybil attacks (one user registers as multiple, seemingly independent users)

Detour: Secure Multiparty Computation (MPC)

- Subfield of cryptography
- Multiple people jointly compute a function on inputs provided by everyone *while keeping those inputs private from each other*
- Adversarial models:
  - Semi-honest / honest-but-curious: Assume participants follow the protocol, but want to learn the private values
  - Malicious: Assume participants may cheat
- Interesting cryptography (you can learn about in other courses)

Simple MPC Example

- Three people want to compute the average # of succulent plants they have without (shamefully) admitting how many they have
- Each person creates three shares of their own count, distributing them securely to the three people (including themselves)
  - Blase has 150. Makes “shares” 120, -20, 50 (sum to 150)
  - Alison has 20. Makes “shares” 100, -120, 40 (sum to 20)
  - Madison has 10. Makes “shares” 50, -60, 20 (sum to 10)
- Blase ends up with (120, 100, 50); Alison with (-20, -120, -60); Madison with (50, 40, 20).
- Averages: Blase (90); Alison (-66.7); Madison (36.7) = **60**

Private Set Intersection (PSI)

- Enables two parties to find the intersection of sets each hold without learning about the other's items not in the intersection
- **Insecure** example: what if we sent each other hashes of everything in our respective sets? Any matches are probabilistically in the intersection (w/ a chance of false positives)
  - We would need to promise not to try hashing elements not in our own sets that we wondered might be in each other's
  - Often assumes "honest but curious" adversary who follows the protocol

Private Set Intersection (PSI)

- Used, for instance, in Apple's password breach monitoring service

The underlying protocol partitions the list of curated passwords, which contained approximately 1.5 billion passwords at the time of this writing, into 2^{15} different buckets. The bucket a password belongs to is based on the first 15 bits of the SHA256 hash value of the password. Additionally, each leaked password, pw, is associated with an elliptic curve point on the NIST P256 curve: $P_{pw} = \alpha \cdot H_{SWU}(pw)$, where α is a secret random key known only to Apple and H_{SWU} is a random oracle function that maps passwords to curve points based on the Shallue-van de Woestijne-Ulas method. This transformation is designed to computationally hide the values of passwords and helps prevent revealing newly leaked passwords through Password Monitoring.

To compute the private set intersection, the user's device determines the bucket the user's password belongs to using λ, the 15-bit prefix of SHA256(upw), where upw is one of the user's passwords. The device generates their own random constant, β, and sends the point $P_c = \beta \cdot H_{SWU}(upw)$ to the server, along with a request for the bucket corresponding to λ. Here β hides information about the user's password and limits to λ the information exposed from the password to Apple. Finally, the server takes the point sent by the user's device, computes $\alpha P_c = \alpha\beta \cdot H_{SWU}(upw)$, and returns it, along with the appropriate bucket of points—$B\lambda = \{ P_{pw} \mid \text{SHA256(pw) begins with prefix } \lambda\}$—to the device.

The returned information allows the device to compute $B'_\lambda = \{\beta \cdot P_{pw} \mid P_{pw} \in B_\lambda\}$, and ascertains that the user's password has been leaked if $\alpha P_c \in B'_\lambda$.

Image from https://support.apple.com/guide/security/password-monitoring-sec78e79fc3b/web

Apple’s Proposals For CSAM Detection

- Broader societal questions:
  - Who controls the database of CSAM?
  - Who decides what objects are included?
  - What happens when possible CSAM is detected?

An Additional Tool That Uses Hashing (Often Used For Local Differential Privacy)

Bloom Filters

- Probabilistic data structure for *set membership*
  - False negatives are *impossible*
    - Bloom filter returns “no” → True answer is “no”
  - False positives are *possible*
    - Bloom filter returns “yes” → True answer is probably “yes,” but might be “no” with some probability (that you can calculate)
- Define an array of **m** bits and **k** different hash functions

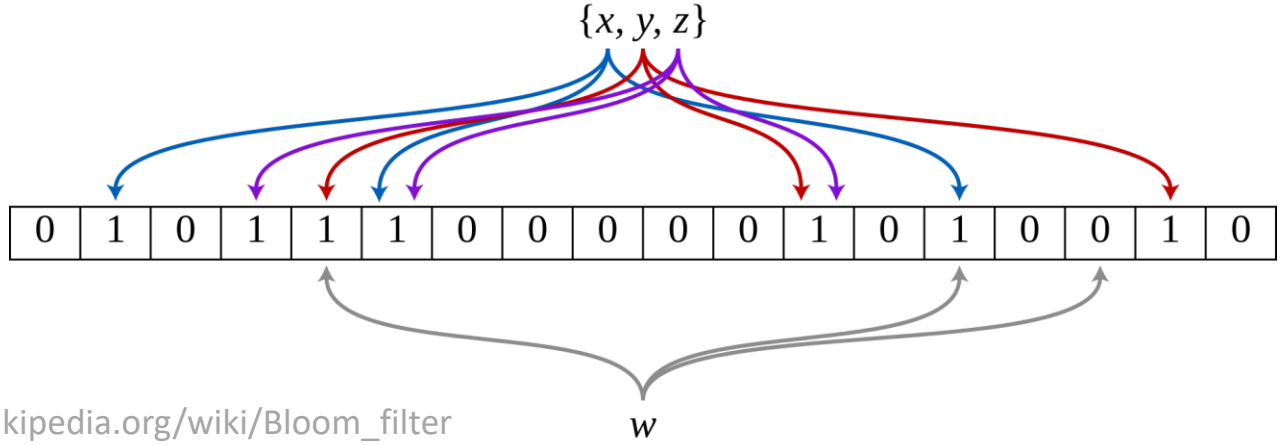

Image taken from https://en.wikipedia.org/wiki/Bloom_filter