

Protecting Visual Assets: Digital Image Counter-Forensics



Nikita Mazurov

nikita.mazurov@mah.se

Kenny Brown

farside792@gmail.com

Why talk about image security?

❖ Increasing prominence of images...

- # of Instagram users jumped from 1 to 500 million from 2010 to 2016.

❖ ...Coupled with increasing image data-mining capabilities

- For example, FindFace claims to be able to link crowd faces to social media profiles with 70% accuracy

❖ ...all combine to form a clear picture indeed: greater vigilance is necessary when developing and deploying image handling strategies

What we'll be doing today

- ❖ Exploring the myriad ways images can be mined for (non-obvious) actionable intelligence
- ❖ Offering up some mitigating counter-forensic & counter-surveillance techniques for image handling
- ❖ Focusing on:
 - **Alteration**
 - **Obfuscation**
 - **Redaction**



What's wrong with this picture?

❖ Take a few moments to jot down all the information you think this image could be leaking

❖ <https://tinyurl.com/justadesk>



A typology of image-handling privacy concerns

- I. Metadata
- II. Secondary Location Leakage & Related Visual Leaks
- III. Safe Redaction Protocol
- IV. Image Discovery (1): Fusking
- V. Image Discovery (2): Content-Based Image Retrieval
- VI. Image Discovery (3): Social Media Mining

I. Metadata

- ❖ Metadata is simply 'data about data'
- ❖ In our case, it is specifically *information about the image* stored in the image file, but not (usually) seen when the image is opened in, e.g., a standard web browser or image viewer
- ❖ Exif (Exchangeable image file format): one (but not the only, e.g. IPTC for keyword tagging) popular standard for image metadata

How to view metadata?

❖ In-browser: Exif Viewer —

<https://addons.mozilla.org/firefox/addon/exif-viewer/>

❖ N.B. Avoid web-based metadata viewer 'solutions' (read: **don't upload MD-laden images anywhere!**)

❖ Stand-alone: ExifTool —

<https://www.sno.phy.queensu.ca/~phil/exiftool/>

Sample photo metadata analysis

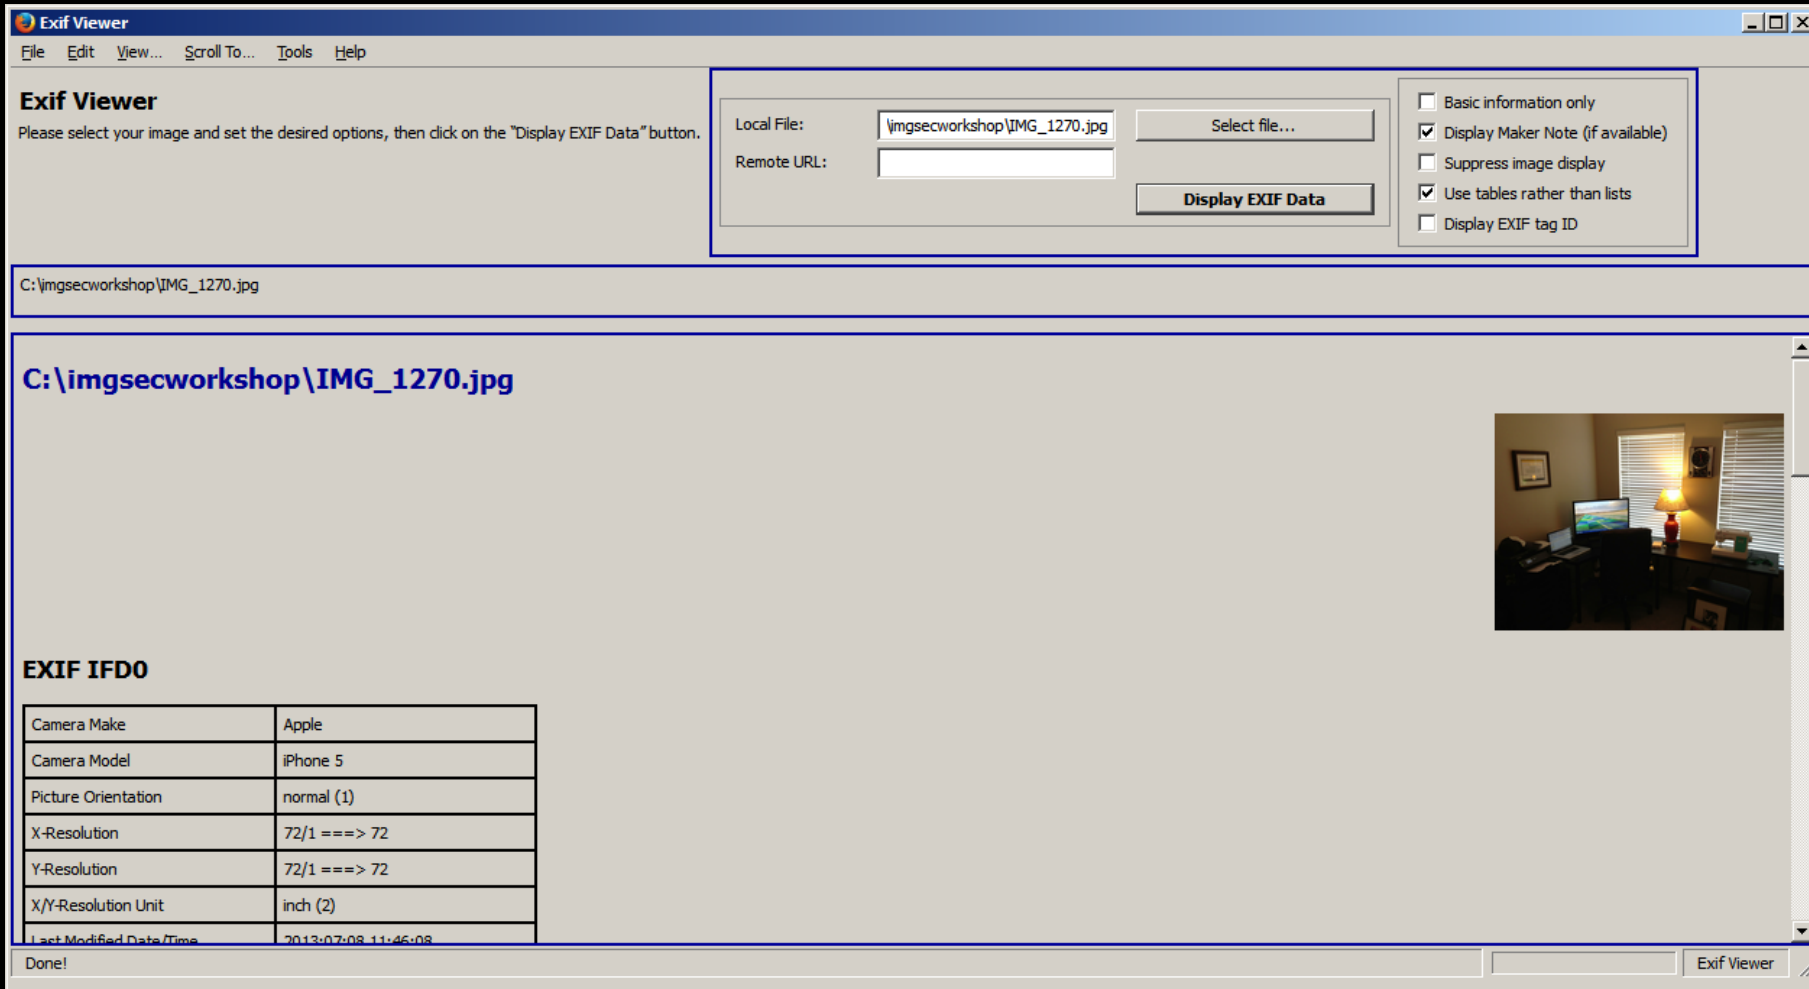
❖ Let's take a look at the metadata in a sample photo file:

- <https://tinyurl.com/insidearoom>



Working with Exif Viewer

❖ (after installing the add-on) Firefox → Tools → Exif Viewer → Select File



The screenshot shows the Exif Viewer application window. The title bar reads "Exif Viewer". The menu bar includes "File", "Edit", "View...", "Scroll To...", "Tools", and "Help". The main area is titled "Exif Viewer" and contains the instruction: "Please select your image and set the desired options, then click on the 'Display EXIF Data' button." Below this, there are input fields for "Local File:" (containing "\imgsecworkshop\IMG_1270.jpg") and "Remote URL:". A "Select file..." button is next to the local file field, and a "Display EXIF Data" button is below the remote URL field. To the right, there are several checkboxes: "Basic information only" (unchecked), "Display Maker Note (if available)" (checked), "Suppress image display" (unchecked), "Use tables rather than lists" (checked), and "Display EXIF tag ID" (unchecked). Below the main area, the file path "C:\imgsecworkshop\IMG_1270.jpg" is displayed. A preview of the image is shown on the right side of the main area. Below the preview, the EXIF data is displayed in a table format under the heading "EXIF IFD0".

EXIF IFD0	
Camera Make	Apple
Camera Model	iPhone 5
Picture Orientation	normal (1)
X-Resolution	72/1 ==> 72
Y-Resolution	72/1 ==> 72
X/Y-Resolution Unit	inch (2)
Last Modified Date/Time	2013:07:08 11:46:08

Done!

Working with ExifTool

❖ (from the command line — or via `exiftool(-k)`):

`exiftool filename.jpg`

(e.g.: `exiftool IMG_1270.jpg`)

Further details:

<https://www.sno.phy.queensu.ca/~phil/exiftool/index.html#running>

```
C:\imgsecworkshop>exiftool IMG_1270.jpg
ExifTool Version Number      : 9.31
File Name                    : IMG_1270.jpg
Directory                   : .
File Size                    : 1622 kB
File Modification Date/Time  : 2016:04:05 02:12:28+02:00
File Access Date/Time       : 2016:04:05 03:03:59+02:00
File Creation Date/Time     : 2016:04:05 03:00:05+02:00
File Permissions             : rw-rw-rw-
File Type                    : JPEG
MIME Type                    : image/jpeg
Exif Byte Order              : Big-endian (Motorola, MM)
Make                         : Apple
Camera Model Name            : iPhone 5
Orientation                  : Horizontal (normal)
X Resolution                  : 72
Y Resolution                  : 72
Resolution Unit              : inches
Modify Date                  : 2013:07:08 11:46:08
Host Computer                 : iPhone (iPhone OS 6.1.4)
Cb Cr Positioning            : Centered
Exposure Time                : 1/24
F Number                     : 2.4
Exposure Program             : Program AE
ISO                           : 50
Exif Version                 : 0221
Date/Time Original           : 2013:07:08 11:46:08
Create Date                  : 2013:07:08 11:46:08
Components Configuration     : -, -, -, Y
Shutter Speed Value          : 1/24
Aperture Value               : 2.4
Brightness Value             : 3.156382079
Metering Mode                : Spot
Flash                        : Off, Did not fire
Focal Length                  : 4.1 mm
```

What information can metadata reveal?

- ❖ Amidst a barrage of photo-technical minutiae (e.g. shutter speed, aperture, brightness, exposure, focal length, et al.), a number of more immediately *actionable elements* may also be present:
 - Camera Make/Model/Serial Number
 - Date/Time/Timezone in which the photo was taken
 - GPS coordinates at which the photo was taken
 - Name of the camera owner
 - Thumbnail of the original image

- ❖ Let's not be too quick to dismiss all that minutiae though!
 - Device fingerprinting: device-specific technical settings could identify source device across disparate photo datasets

Camera data (make/model/serial)

- ❖ The camera make (i.e. brand) can be used to link a photographer's camera to a photograph
- ❖ The specific model of a given make or brand can further strengthen the link
- ❖ A specific serial number can then serve as the final, fatal tightening of the noose

Serial numbers: special consideration

❖ If the MD contains a serial number, questions adversaries may ask include:

- Was a product registration form filled out upon purchase of the device?
If yes → manufacturer may be able to provide owner info (name, address, phone, email)
- Was the device included in an insurance inventory/asset list?
If yes → insurance provider will be able to supply the aforementioned data (may already be preemptively working with State Actors (e.g. Local/National/Foreign Law Enforcement))
- Are there any other photos online which have the same serial?
If yes → do any of the other photos (*or the websites on which they're hosted*) reveal any actionable intelligence?
 - <http://cameratrace.com/>
 - <http://www.stolencamerafinder.com/>
 - (as well as just a Google search for the SN)

Date and time

- ❖ A photo may contain various unique timestamps, including:
 - Date the photo was taken
 - Date the photo was last modified
 - Date the GPS coordinates were recorded
- ❖ If the photo is either known/suspected to be taken at a given location, CCTV footage can be reviewed for the corresponding date/time to streamline subject identification (subject can then be tracked across various CCTV vectors to, e.g., a given vehicle or office)
- ❖ The time zone may narrow down the location at which the photo was taken (as well as corroborating GPS data, if available)

GPS data

❖ Global Positioning System (GPS) coordinates are accurate within ~3-10 meters (~10-33 feet)

- E.g. “Results indicate that A-GPS locations obtained using **the 3G iPhone** are much less accurate than those from regular autonomous GPS units (average **median error of 8 m** for ten 20-minute field tests) but appear sufficient for most Location Based Services (LBS)”
 - (Zandbergen, P. A. (2009), “Accuracy of iPhone Locations: A Comparison of Assisted GPS, WiFi and Cellular Positioning”. *Transactions in GIS*, 13: 5–25. doi:10.1111/j.1467-9671.2009.01152.x).
 - Can assume it's more accurate now
- i.e. probably good enough to tell what building or immediate vicinity the photo was taken in

Camera owner name

- ❖ Some cameras (and other devices) encourage you to customize your camera upon purchase, by adding your name during the initial setup procedure

DON'T

- ❖ Corollary: similarly, don't assign identifying names to SD cards or folders (avoid custom, potentially-compromising volume and directory naming)

Thumbnail data

- ❖ Exif MD is not necessarily confined to *text*, but can contain *binary* data (e.g., an imbedded thumbnail image)
- ❖ Say a photo is cropped in a photo-editing suite to remove compromising components of the image
 - The original, uncropped photo may still be in the imbedded thumbnail
- ❖ Windows also includes its own separate, hidden thumbnails database file (thumbs.db) in image folders by default
 - Instructions for disabling: <http://www.pcworld.com/article/2999243/windows/manage-thumbs-db-files-in-windows-and-on-the-network.html>

Metadata deletion

- ❖ Finally some good news: it is very, very easy to delete photo MD; requiring only one command:

```
exiftool filename.jpg -overwrite_original -all=
```

- ❖ Can also scrub entire directories, e.g.:

```
exiftool c:\photostoclean\ -overwrite_original -all=
```

- ❖ Or by dragging a file or folder onto a copy of Exiftool that has been named as: `exiftool(exiftool -overwrite_original -all=).exe`

Sorry, metadata deletion might not actually be that easy :(

- ❖ Some metadata may be termed *persistent*, or harder to delete.
- ❖ Exiftool may have trouble automatically wiping some MD
 - E.g., PNG text chunks: iTXt, tEXt, & zTXt
- ❖ Open the image in a hex editor to make sure MD fields have all been wiped; conduct manual wiping if necessary
- ❖ Highlights the dangers of over-reliance on automation

Best practice: deletion-by-default

❖ Default operations protocol should be: Delete all MD *unless you have a good reason to keep it*;

- NOT 'keep all MD unless you have a good reason to delete it'

❖ MD from seemingly innocuous images may be used to glean intelligence from MD-removed sensitive images

- Consider three images A, B, C taken on a trip. A and C deemed harmless, B has MD wiped. Location of B may be estimated based on MD of A & B.

❖ Vendors: implement deletion-by-default into image-handling workflows (with warnings for disabling)

Metadata modification

- ❖ Instead of outright deletion, modification may be desirable.
- ❖ Many Exif values can be changed to one's liking (as long as one knows the proper, at times non-intuitive, value (or tag) name).
- ❖ <https://www.sno.phy.queensu.ca/~phil/exiftool/TagNames/>

MD modification case study: spoofing GPS coordinates

- ❖ Let's change the coordinates of our sample IMG_1270.jpg from Manisa to Vegas.
- ❖ First step: find the desired locale's coordinates.
 - <https://maps.google.com>
 - _ Put in the desired location name → right-click and select 'What's here' → click on the decimal degree coordinates → get the degree/minute/seconds format
 - (could also use <https://www.openstreetmap.org> , though would need to use a third-party to do decimal-to-degree conversions)

Sample Google Maps coordinate discovery procedure

The image shows a Google Maps interface with several key elements and annotations:

- Search Bar:** Contains the coordinates `36.088446, -115.177843`.
- Coordinate Display:** Shows `36°05'18.4"N 115°10'40.2"W` and `36.088446, -115.177843`. A red box highlights the coordinates, with a red arrow labeled **3** pointing to it.
- Map:** Displays a street view of Las Vegas, NV, with a red pin and a context menu. A red arrow labeled **1** points to the **What's here?** option in the menu.
- Context Menu:** Includes options like "Directions from here", "Directions to here", "What's here?", "Search nearby", "Print", "Add a missing place", "Report a data problem", and "Measure distance".
- Location Card:** Shows "Mandalay Bay" with address "3950 S Las Vegas Blvd, Las Vegas..." and coordinates `36.088481, -115.177832`. A red arrow labeled **2** points to this card.
- Bottom Panel:** Features buttons for "SAVE", "NEARBY", "SEND TO YOUR PHONE", and "SHARE". Below these are options to "Add a missing place" and "Add a label".


Metadata injection with ExifTool

```
exiftool IMG_1270.jpg -GPSLatitude="36 deg 05', 18.4'''' -  
GPSLongitude="115 deg 10', 40.2'''' -GPSLongitudeRef=W -  
overwrite_original
```

❖ (can omit `-overwrite_original` during testing)

- Additional GPS tags:

<https://www.sno.phy.queensu.ca/~phil/exiftool/TagNames/GPS.html>






Exercise extreme caution: modification is trickier than deletion

- ❖ MD modification requires more care than MD deletion
- ❖ When spoofing MD, avoid potential future forensic detection of MD tampering by paying careful attention to **MakerNote specifications**.
- ❖ What are MakerNotes? Manufacturer-specific metadata tags (e.g. images taken with Sony devices may have some MD fields not present in Olympus-sourced images)
- ❖ Refer to manufacturer-specific entries on <https://www.sno.phy.queensu.ca/~phil/exiftool/TagNames/> for explanatory lists of MakerNote tags

Assume *adversarial familiarity* with the given terrain

- ❖ Effective counter-forensic threat modeling: not underestimating the extent of an adversary's familiarity with given the subject field




The Exif standard also supports data called “maker-notes.” These data fields and their values are unique to each digital camera manufacturer. They can help determine if a suspect has tampered with Exif data in an attempt to prevent linking images to a specific digital camera. For example, encountering an image with the Exif data of a Canon camera and the makernotes of a Nikon would indicate that fields have been modified.

II. Secondary Location Leakage & Related Environmental/Visual Leaks

- ❖ Be cognizant of all manner of visual clues (both latent and apparent) which may inadvertently compromise the situational security of the image. Including, but not limited to, the following **localisms**:
 - Brand names
 - Native flora/fauna
 - Any textual data (e.g. signage, newspapers)
 - Light switches
 - Electrical outlets
 - Identifiable locations (e.g. landmarks, chain hotel rooms)

Electrical outlet geolocation



Additional visual clues

- ❖ Aside from *localisms* (which betray the locality an image was taken in), be wary of personally identifiable *slippage*, such as:
 - Reflections
 - Exposed body parts
 - Height revelations (e.g. if standing next to a measurable object)
 - Location calculations (e.g. if photographing an event, can the location of the photographer be deduced by forensic analysts afterwards via angle-measurement and CCTV footage?)
 - All boils down to: (any) **extraneity** (in a photo) = **vulnerability**

III. Safe Redaction Protocol


- ❖ Potentially sensitive components of images should be redacted, not blurred
- ❖ Selective brightness of blurred image components can be reversed (e.g. a blurred '5' will look different than a blurred '1'; by testing blur algorithms on various font-sets, may be possible to deduce the original text)
 - Dheera Venkatraman, "Why blurring sensitive information is a bad idea"
<https://dheera.net/projects/blur>
- ❖ (Don't swirl or use other gimmicky, novelty effects either)





Remember: everyone matters in a photo

- ❖ Incomplete redaction can lead to deanonymization of the redacted component
- ❖ Example: a photo of a human subject redacts the human, while a canine co-traveler is not redacted. Canine is then linked to the human via, e.g.:
 - pet store purchasing records
 - veterinary records
 - other service records (grooming)
 - local dog park and neighborhood surveillance



Overdeletion is preferable to underdeletion

- ❖ Redact more than necessary
- ❖ Avoid leaking remainder information
 - E.g., make sure there are no ascender/descender remnants
- ❖ Redact empty space to foil probable word-size attacks

2006 Cat 980G Wheel Loader, 2200 hours, runs well c/w 2 buckets, grapple, chains, manual. \$120,000 OBO. [Redacted] Langley, BC.

2006 Cat 980G Wheel Loader, 2200 hours, runs well c/w 2 buckets, grapple, chains, manual. \$120,000 OBO. [Redacted]

2006 Cat 980G Wheel Loader, 2200 hours, runs well c/w 2 buckets, grapple, chains, manual. \$120,000 OBO. [Redacted]

IV. Image Discovery (1): Fusking

- ❖ 'Fusking' is the exploitation of the practice of camera manufacturers to name images sequentially to find images which one may not wish to be seen.
- ❖ For example, if you give someone a link to http://yoursite.com/images/IMG_0001.jpg, could they simply scrape the directory for IMG_[0000-9999].jpg?

Common fudging patterns

❖ Common photograph prefixes include:

- IMG_####.jpg
- DSC_#####.jpg
- DSCN####.jpg
 - _ also reveals a Nikon camera was used: Digital Still Capture Nikon

❖ But not just limited to numerical sequences:

- Can launch dictionary attacks for common names (e.g. 'vacation.jpg'; 'kids.jpg', etc...)

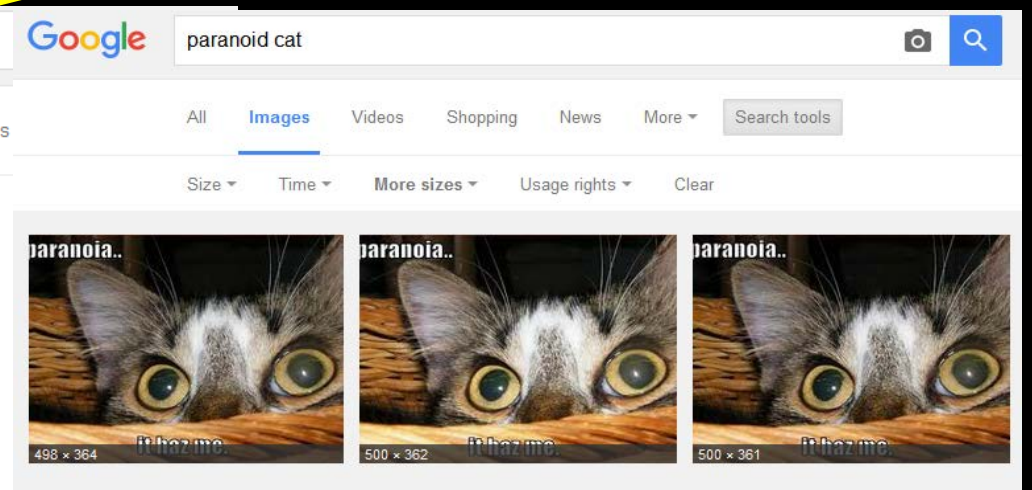
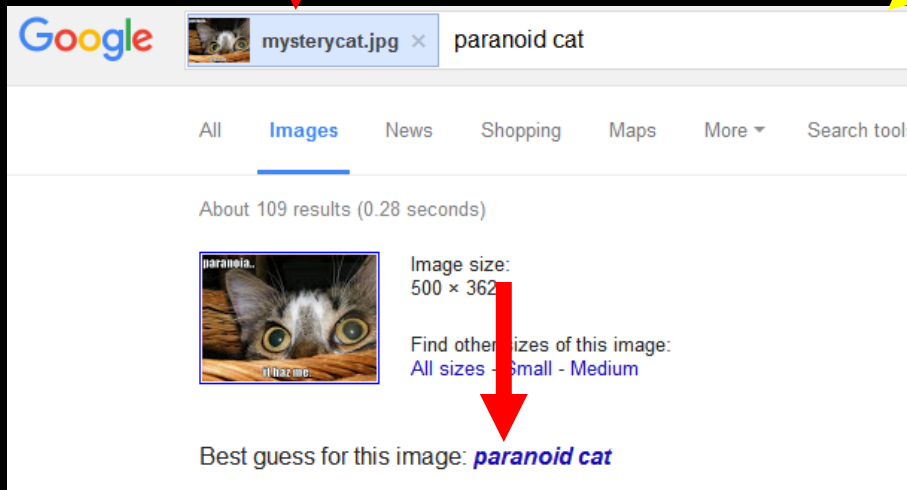
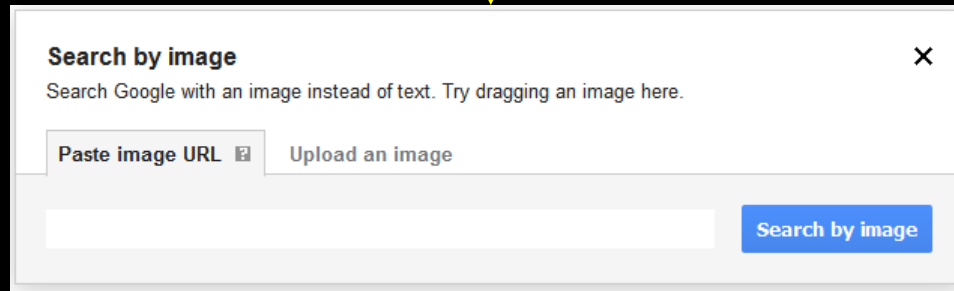
V. Image Discovery (2): Content-Based Image Retrieval

- ❖ CBIR systems search for images based on image *contents*, as opposed to image metadata (e.g. by searching for images which *look like* they have cats (e.g. have a 'cat-like' shape), versus images which are named 'cat' or have been tagged with the keyword 'cat')
- ❖ One common way CBIR search systems can operate is via ***reverse image searching***: querying a search engine *by image* instead of *by keyword*.
 - Instead of searching Google Images by typing 'cat', we can search Google Images by uploading a picture of a cat to find other pictures of cats, ***or to find pictures of the same cat***.



mysterycat.jpg

<https://images.google.com/>



CBIR security considerations

- ❖ Image components may be isolated from a composite image to facilitate, e.g, individual or landmark identification
- ❖ Still frames (screenshots) from video can likewise be used as search queries
- ❖ Run the entire image/cropped selections thereof through reverse image searches preemptively

VI. Image Discovery (3): Social Media Mining

- ❖ Once a SM account is discovered (via, e.g., CBIR), can be utilized for acquaintance mapping
 - E.g., if CBIR leads to a secondary (acquaintance) SM account, can be escalated to in turn find the name of the target's SM account
- ❖ Images and image captions can then further be leveraged for intelligence gathering
 - A photo showing a birthday celebration with the subject wearing a 'birthday girl' hat can be matched to the date posted, to obtain subject's DOB
 - Recently posted location photos reveal subject's immediate location (e.g., if at restaurant, home and car likely both empty and vulnerable)

Vendor responsibility

- ❖ Secure image-handling should not be all on the end-user's shoulders
- ❖ Vendors who deal in products that involve image-handling should implement image sanitization into the product
 - Integrate user privacy into the product workflow from the ground up
- ❖ Fail-safe defaults, with warning screens for potentially unsafe toggles
- ❖ If dealing with cloud-based services, minimize liability by minimizing data retention

Preliminary case study, redux



- ❖ Returning now to the sample image we looked at during the beginning of our session: <https://tinyurl.com/justadesk>
- ❖ What information did you previously jot down? What information can you now extract from it?

Black Hat sound bites: key takeaways

Visual information leakage may be **non-obvious**; therefore...

- ❖ Always remove (**alter, obfuscate, redact**) as much information as you can, even if it's seemingly innocuous
- ❖ Be weary of not just technical leaks (e.g. metadata), but of environmental leaks (e.g. wall sockets)
- ❖ Keep in mind the broader ecosystems your image may propagate in (e.g. friends' social media feeds)

And finally...

“Whenever there’s any doubt, there is no doubt”

Questions?



Comments?

Nikita Mazurov

nikita.mazurov@mah.se

Thank you!

Kenny Brown

farside792@gmail.com