A Picture is Worth a Thousand Words, Literally

Deep Neural Networks for Social Stego

ZEROFOX®

802secure

Philip Tully | Mike Raggo

DC25: Community, Discovery and the Unintended Uses of Technology

2600: The Hacker Quarterly

25 years

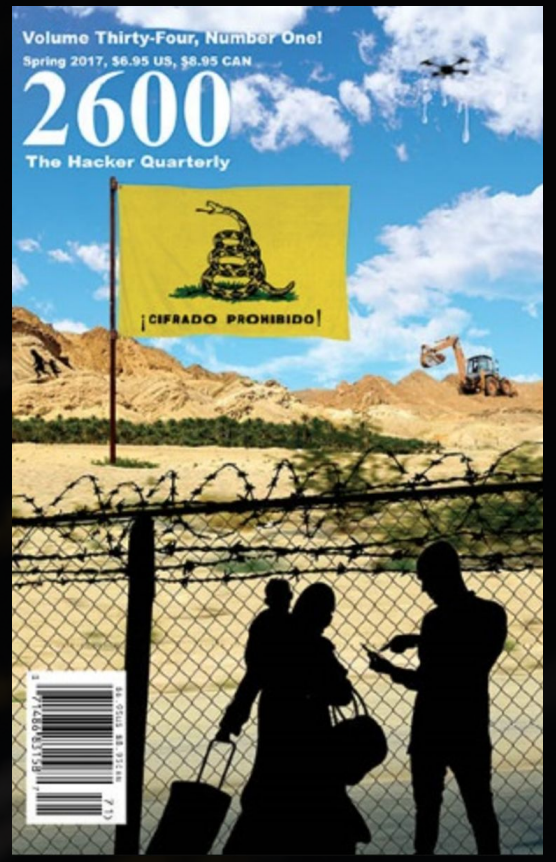

Summer 1992

Summer 2017

#whoami

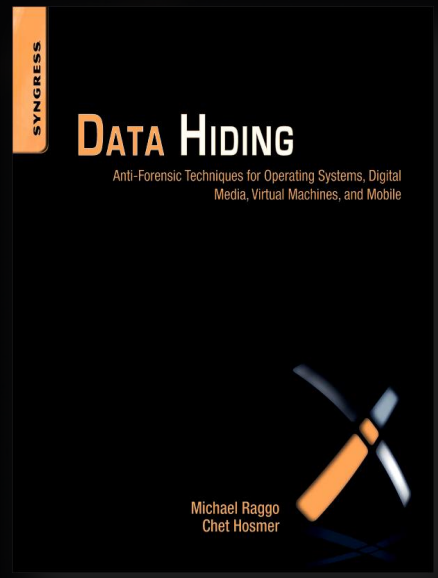

Philip Tully
@phtully

Mike Raggo
@datahiding

ZEROFOX®

802secure

Principal Data Scientist at ZeroFOX

PhD (KTH & University of Edinburgh)

Machine Learning and Neural Nets

CSO @802 Secure, 17 yrs Stego Research

StegSpy DC12, Author “Data Hiding”

NSA National Cryptologic Museum

The Evolution of Steganography

DIY Social Steganography

Deep Neural Networks for Social Stego

Data-Driven Red and Blue Teaming

Wrap Up

A Picture is Worth A Thousand Words:

Deep Neural Networks for Social Stego

The Evolution of Steganography

A Picture is Worth A Thousand Words:

Deep Neural Networks for Social Stego

Covert Communication

“. . . any communication channel that can be exploited by a process to transfer information in a manner that violates the system's security policy.”

Source: U.S. Department of Defense. Trusted Computer System Evaluation “The Orange Book”. Publication DoD 5200.28-STD. Washington: GPO 1985

Evolution of Methods

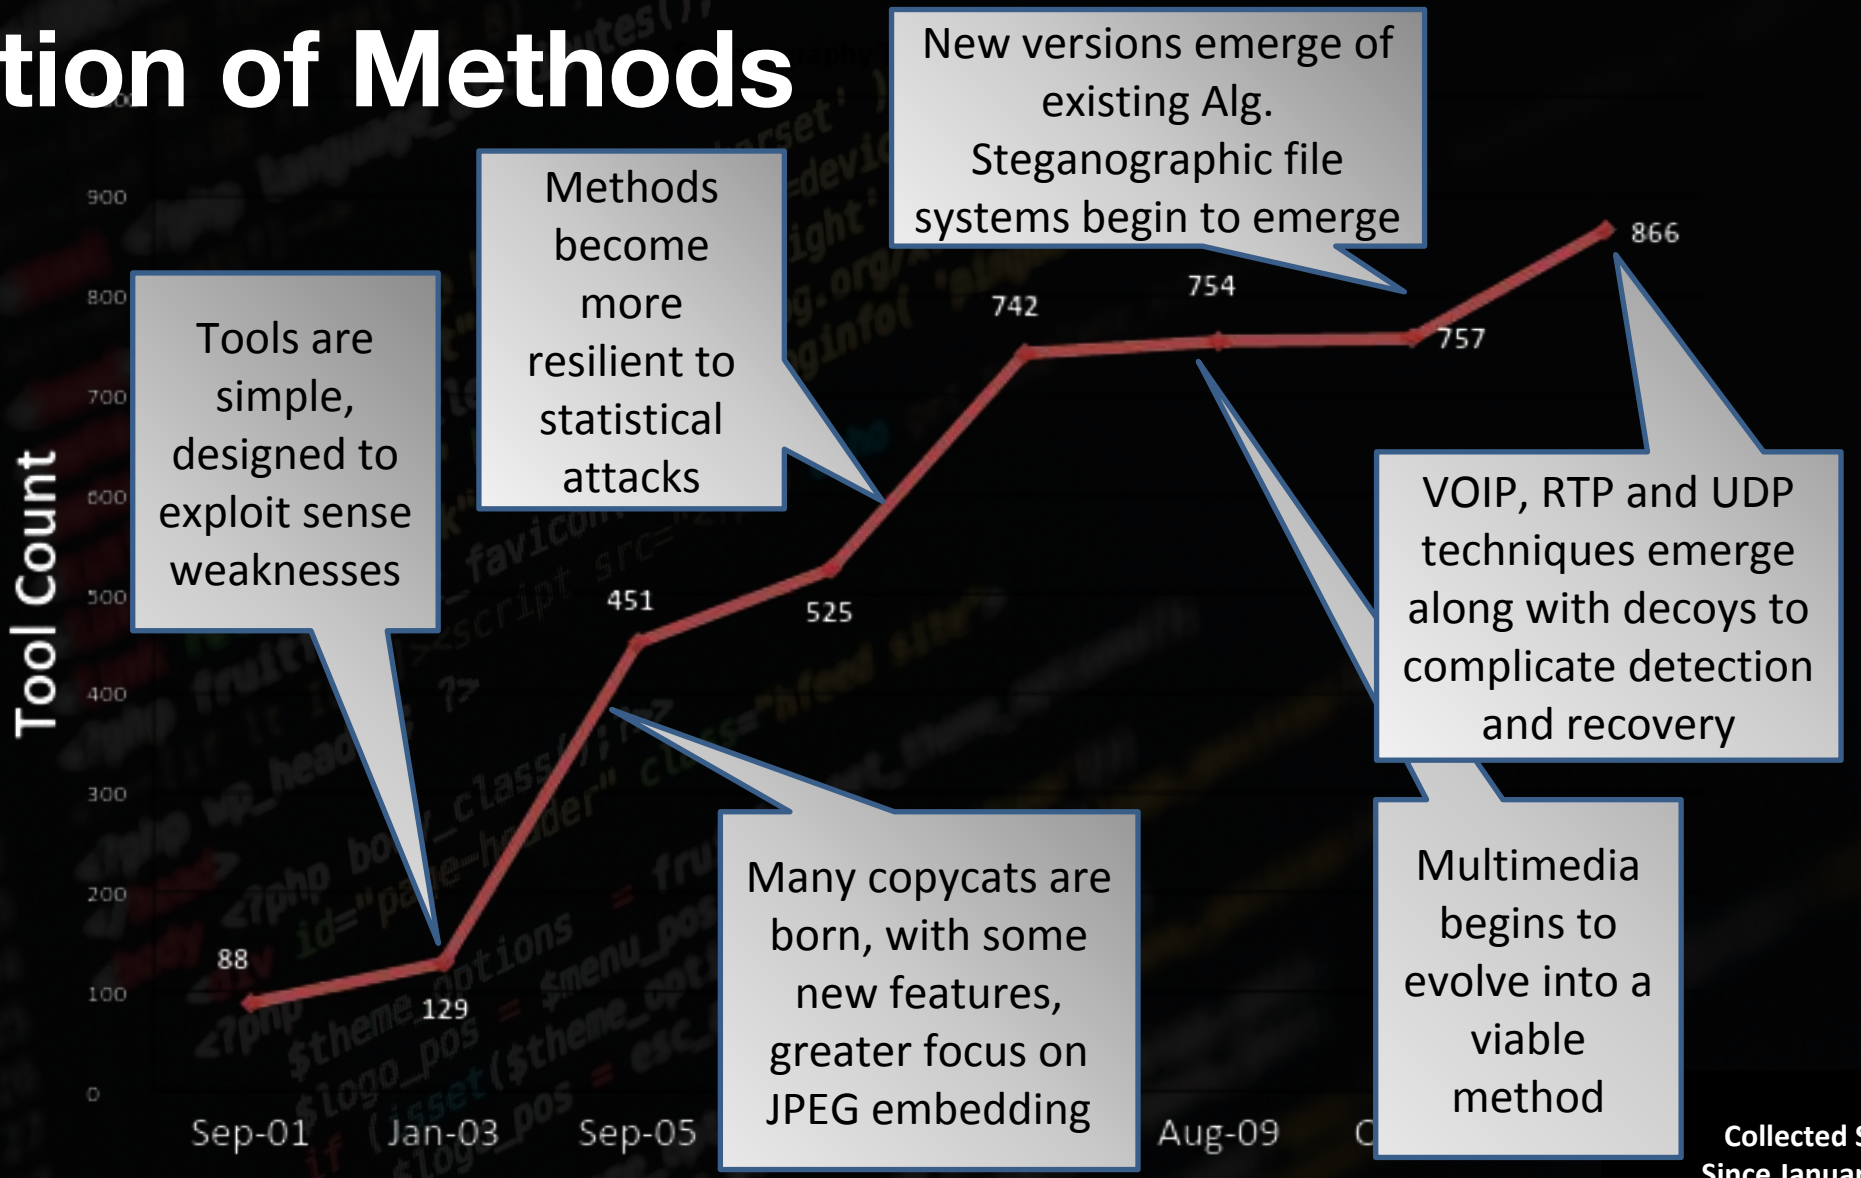

WetStone Labs
Collected Steganography Programs
Since January 1999 Includes versions

Evolution of Stego in the Internet Era

- Stego Apps Decoy Techniques (OpenPuff)
- Stealth Alternate Data Streams (NT)
- Weaponized CnC - Operation Shady RAT (McAfee)
- Prootocols - VOIP, RTP, UDP => WiFi StegoStuffing, Bluetooth (Hosmer/Raggo - Wall of Sheep/Skytalks DC23 & 24)
- MP3 ID3 Metadata exploitation - Hosmer/Raggo Skytalks DC24
- SmartWatch SWATtackhide.py Tizen SDK - Mike Raggo - DEF CON 24 Demo Labs and Wall of Sheep

Types of Steganography

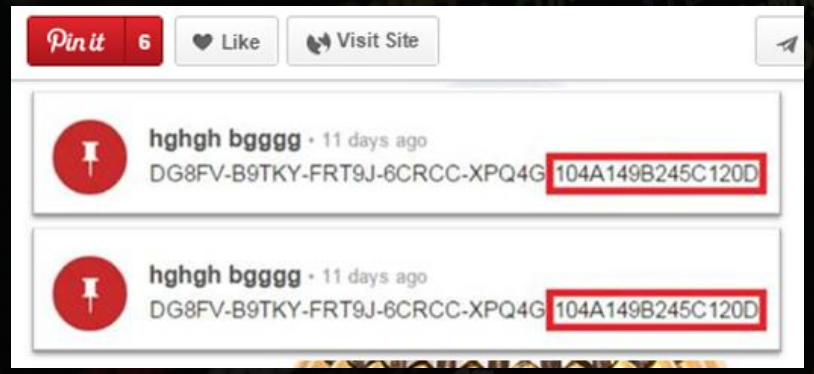

- Text/Linguistic Stego - using Natural Language
- Image
  - Spatial (e.g. LSB)
  - Frequency (DCT/DWT)
  - Metadata (varies by file type and versions) - JPEG EXIF vs. JFIF
- Audio
- Video
- Protocols
- Use of crypto with stego
  - Vigenere, base64, XOR, etc.

DIY Social Steganography

A Picture is Worth A Thousand Words:

Deep Neural Networks for Social Stego

Signals in the Social Noise

4.75 billion
pieces of content shared per day.

80+ million
images uploaded per day.

100+ hours
of video uploaded per minute.

5 billion
+1's per day.

500+ million
tweets per day.

Social Network Image Proliferation

- Image-based social networks have the fastest growing user bases
- Image-based social networks enjoy the highest daily time spent by users
- “Photos or Images” is the content category most frequently share by users
- Social posts containing images produce 650% higher engagement than text alone

Social Stego in the Wild

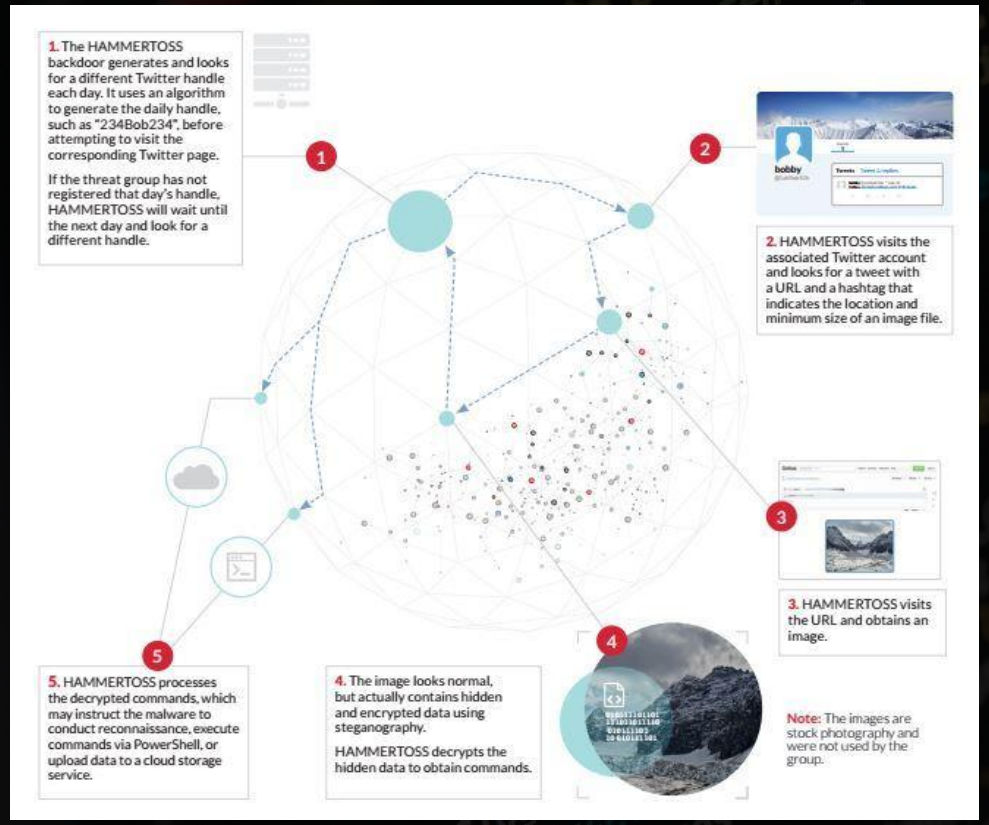

Black Hat: HAMMERTOSS

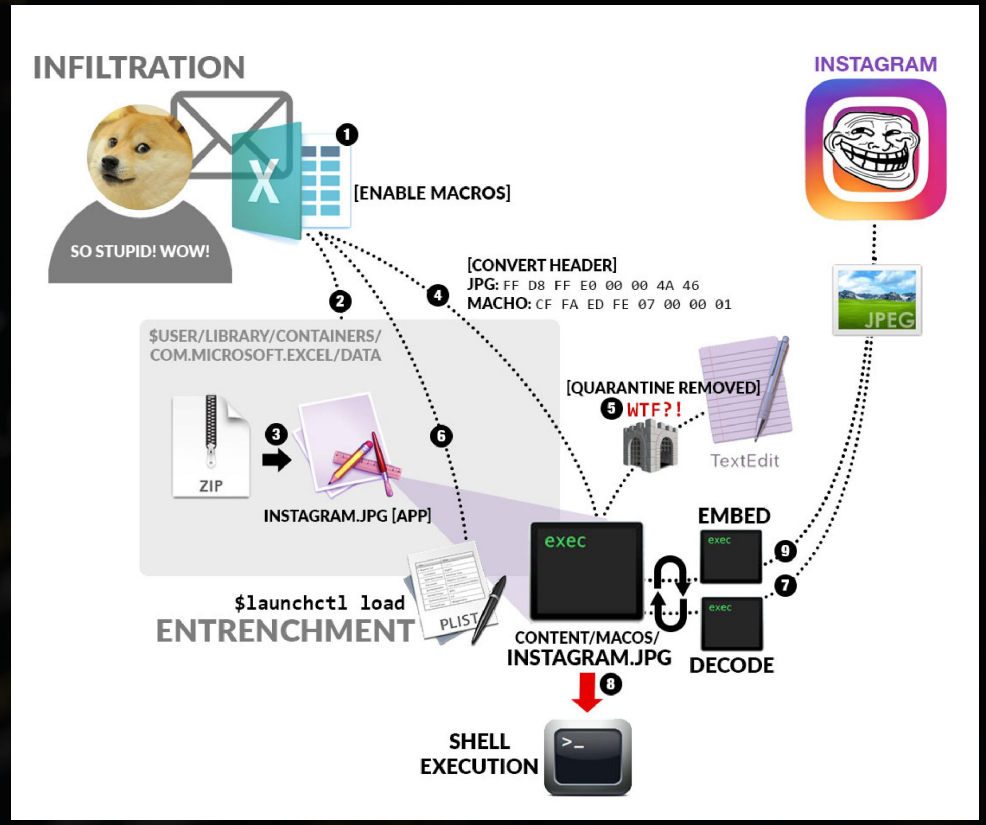

White Hat: Instegogram

Social Network Photo Targets

- Profile Image
- Background Image
- Posted Image(s)
- Photo albums
- DM images

Carrier Image File Types

- Image quality properties:
  - Lossy v. Lossless Raster Compression
  - DPI/PPI
- Common file formats:
  - JPEG (Lossy)
  - PNG (Lossless)
  - TIFF (Lossless)
  - GIF (Lossy)
  - BMP (Lossy)

Trial and Error - Attempted Methods

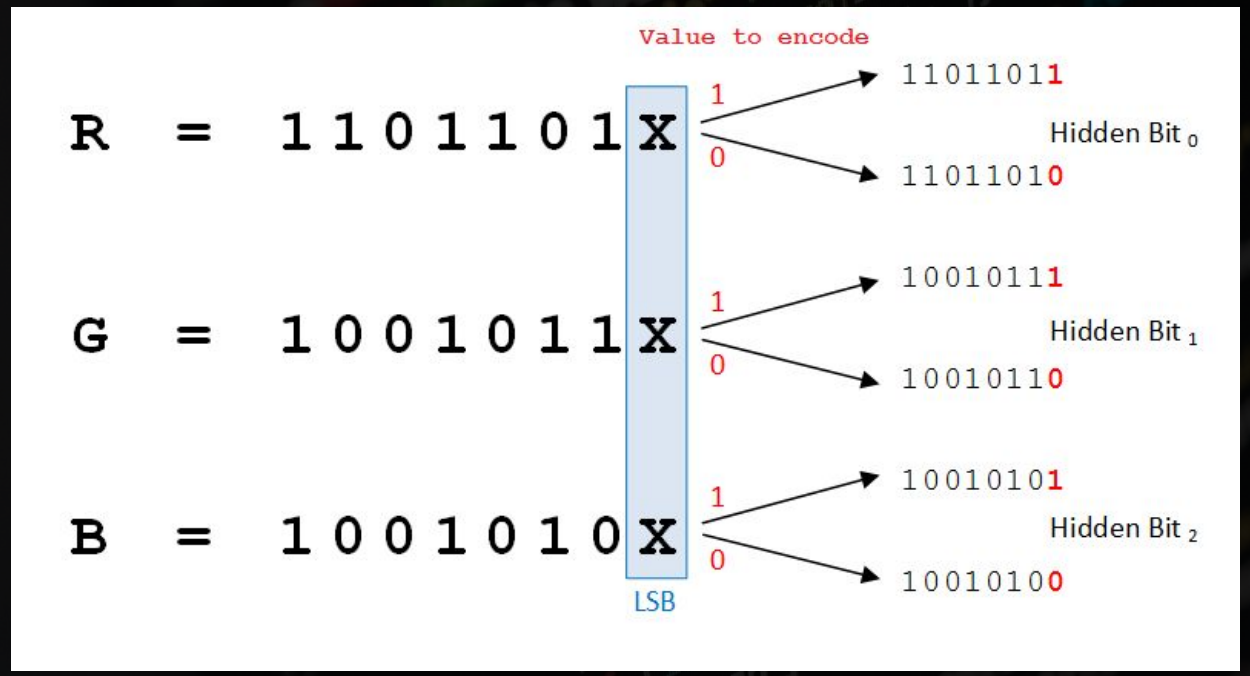

DataGenetics

- Metadata fields (varies by image types JPEG EXIF vs. JFIF, etc.)
- LSB - Least Significant Bit
- Insertion
- Append after EOF marker
- Pre and Post Upload
- Linguistic Steganography

High-Level Testing Workflow

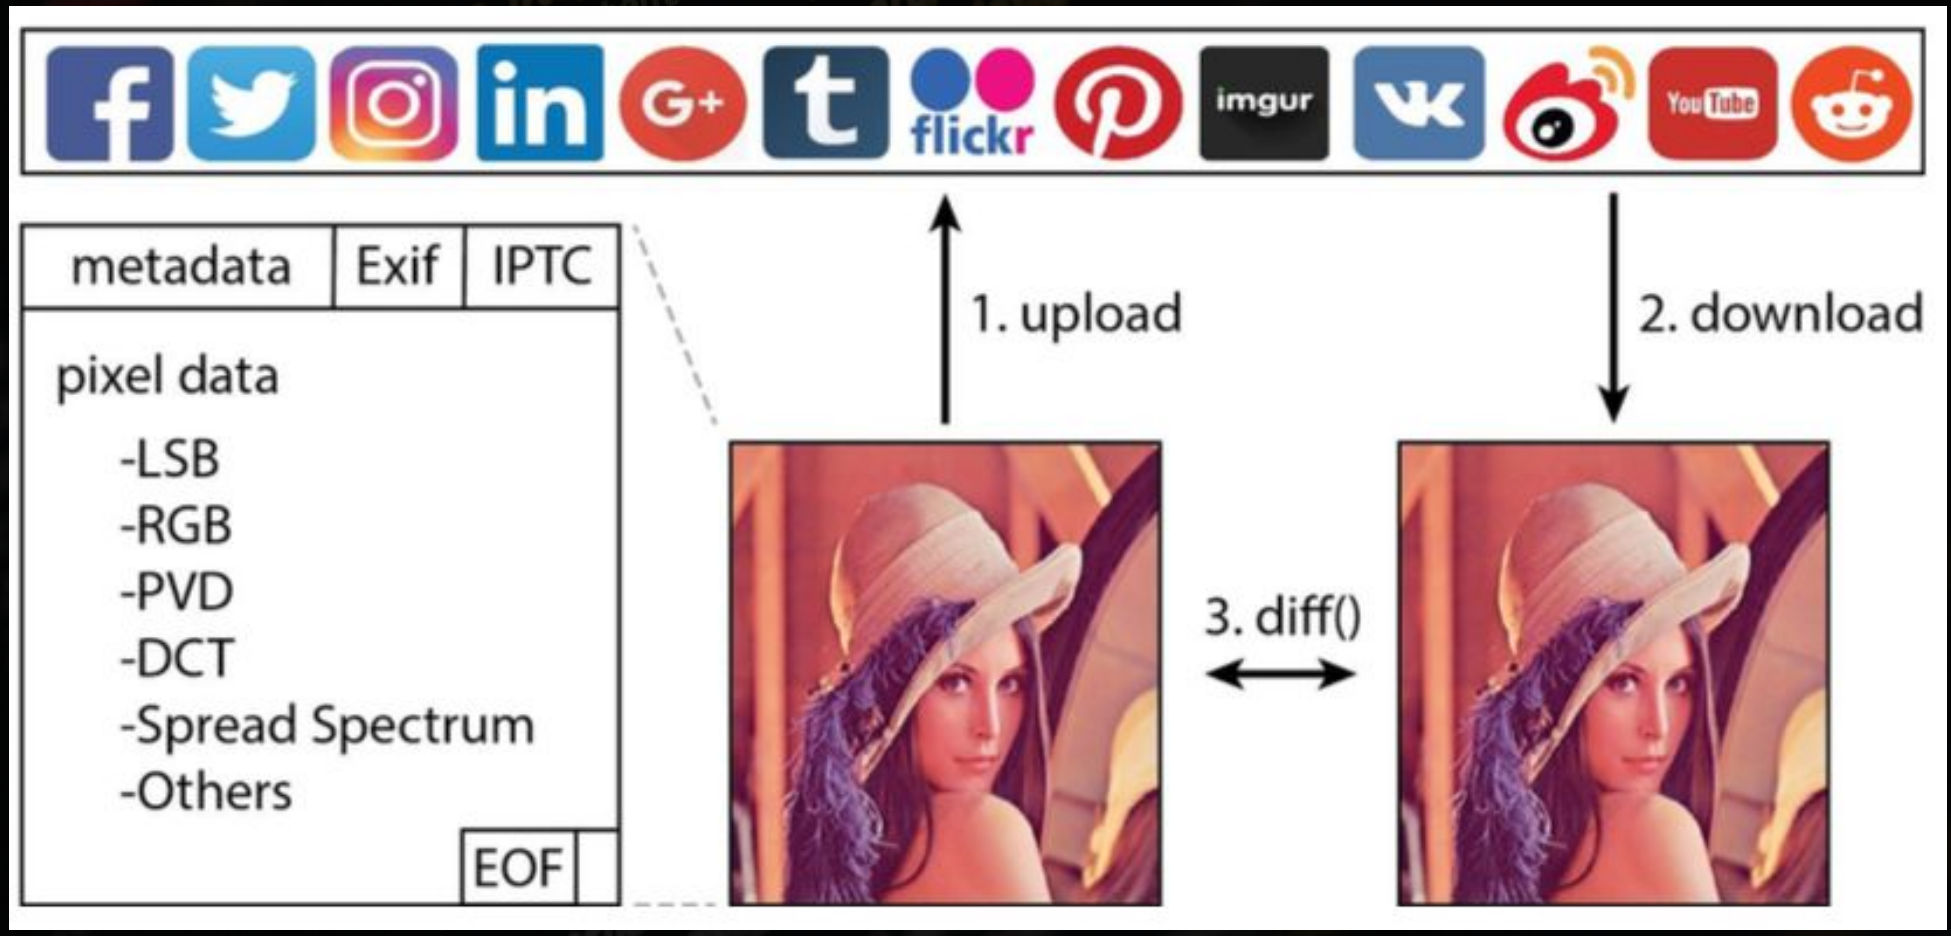

Social Network Data Hiding Survivability

Social Network	Profile Photo	Post an Image	Background Image
Twitter	No	No	No
Facebook	No	No	No
Pinterest	No	Yes	
Instagram	No	No	No
Slack		Yes	
Tumblr	No	No	No
Google+		Yes	

Deep Neural Networks for Social Stego

A Picture is Worth A Thousand Words:

Deep Neural Networks for Social Stego

Secretbook by Owen Campbell-Moore

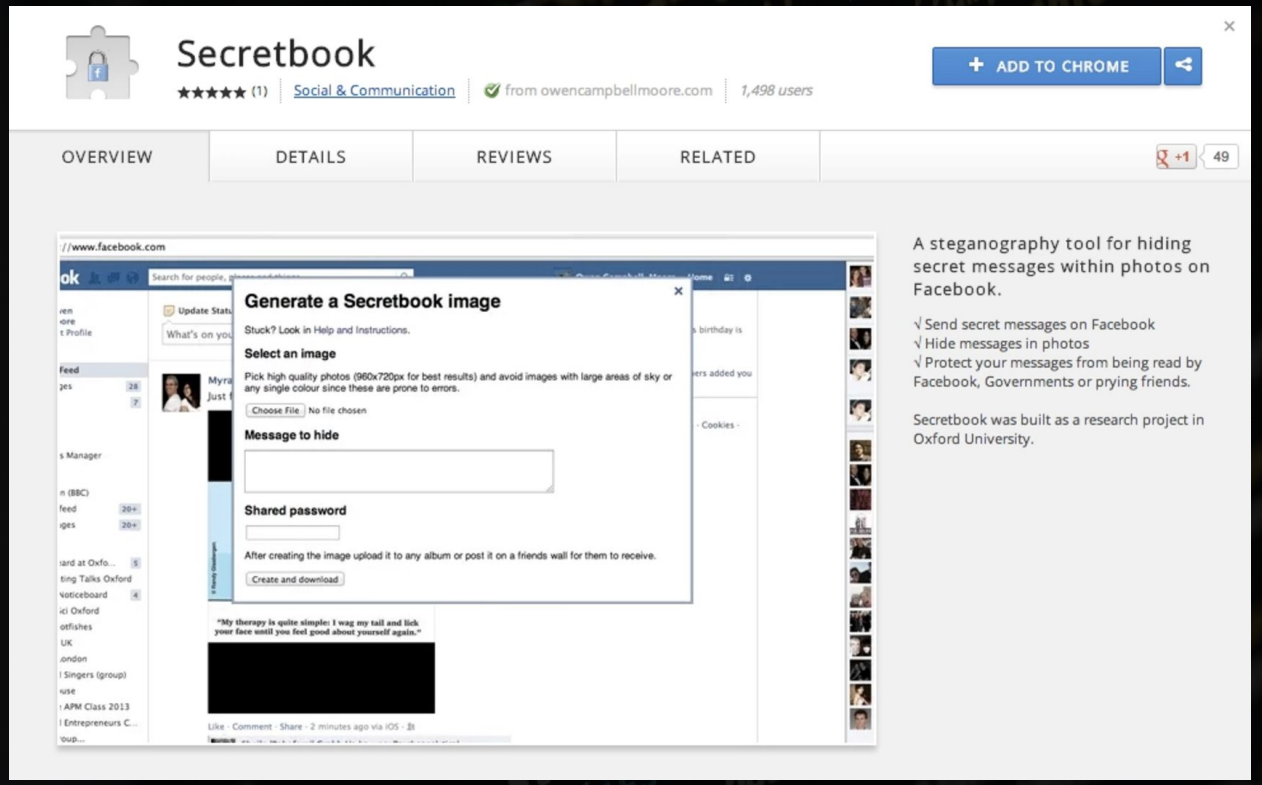

- Open-source Social Stego tool
- Chrome Extension (2013)
- Reverse engineered Facebook's lossy compression algorithm
- Allowed for payloads of up to 140 characters in length

Jamming Techniques

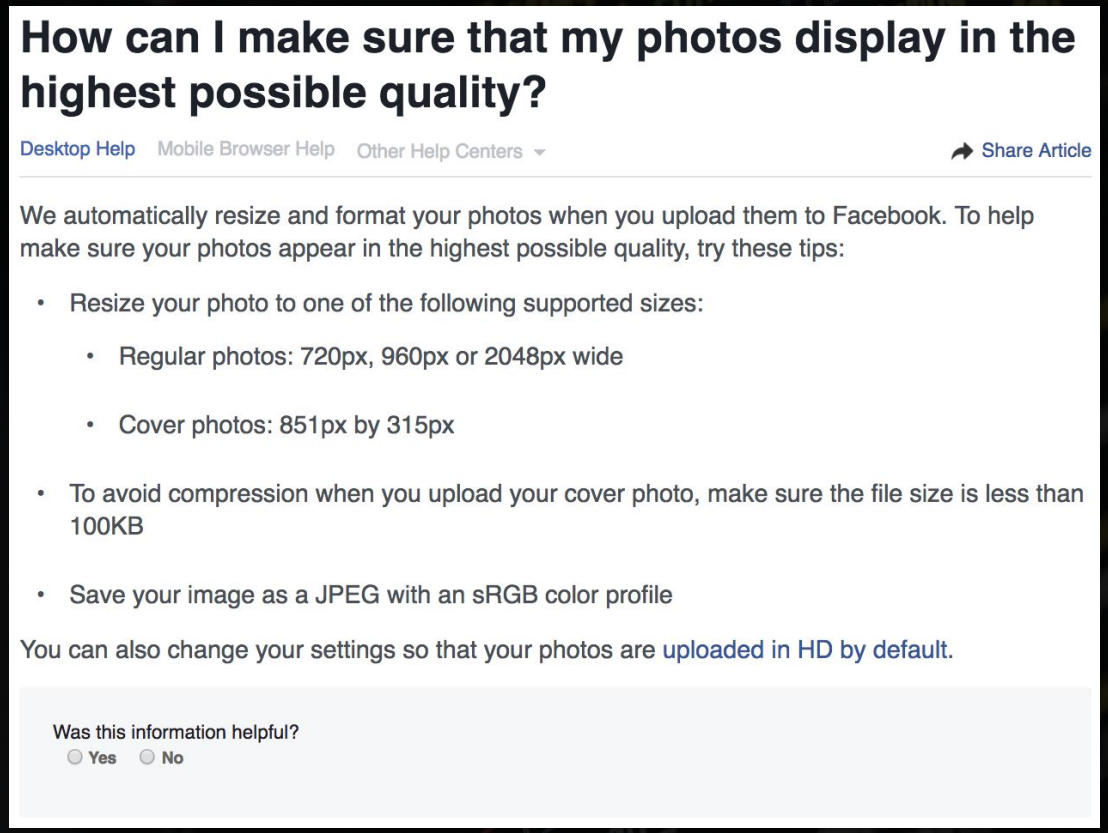

How can I make sure that my photos display in the highest possible quality?

Desktop Help Mobile Browser Help Other Help Centers Share Article

We automatically resize and format your photos when you upload them to Facebook. To help make sure your photos appear in the highest possible quality, try these tips:

- Resize your photo to one of the following supported sizes:
  - Regular photos: 720px, 960px or 2048px wide
  - Cover photos: 851px by 315px
- To avoid compression when you upload your cover photo, make sure the file size is less than 100KB
- Save your image as a JPEG with an sRGB color profile

You can also change your settings so that your photos are uploaded in HD by default.

Was this information helpful?
Yes No

- Server-side image upload restrictions and alterations
- Also legal concerns
  - Crime investigations
  - Trademark infringement
- Common Image upload Alterations:
  - Recompression
  - Metadata stripping
  - Filetype conversion
  - Resizing

Bulk Image Uploads/Downloads

- Data Acquisition made easy
  - Permissive APIs for content creation
  - More content=more engagement=profit
- Off-the-shelf photo aggregators
  - Facebook albums
  - Pinterest boards
  - Flickr sets
  - Google+ Collections
- Or we can do it the ‘hard way’
  - for photo in album{
    upload(photo); sleep(randInt); }

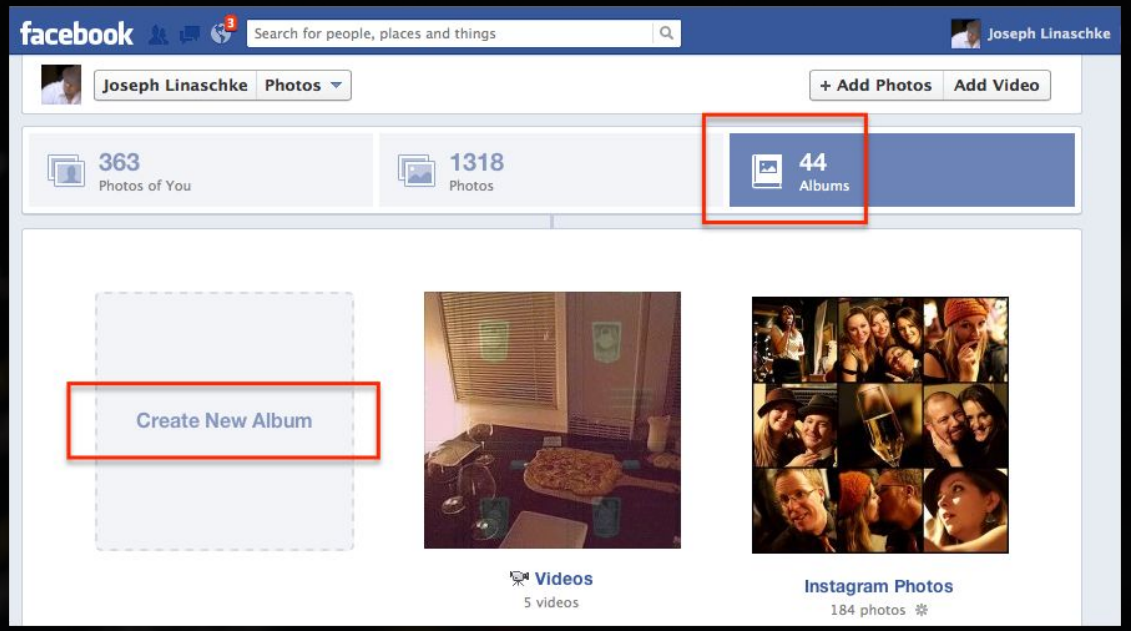

Auto-Generating Data

- Select 50k ImageNet samples
- Automate uploads and downloads
- =100k pre-uploaded and downloaded images
- Compare pixels between phases
- Can comparison/location be automated?
- But Neural Nets don't scale to Images
  - width * height * 3 channels = unmanageable # weights
  - encode these properties into the architecture

What humans see

08 02 22 97 38 15 00 75 04 05 07 78 52
49 49 99 40 17 81 18 57 60 87 17 40 98
81 49 31 73 55 79 14 29 93 71 40 67 53
52 70 95 23 04 60 11 42 69 24 65 56 54
22 31 16 71 51 67 63 89 41 92 36 54 22
24 47 32 60 99 03 45 02 44 75 33 53 78
32 98 01 20 64 23 67 10 26 38 40 67 59
67 26 20 68 02 62 12 20 95 63 94 39 63

What computers see

Convolutional Neural Networks

- Proven great for Computer Vision Tasks:
  - Object classification, Facial recognition
- Pose as a Regression Task
  - Locate optimally embeddable pixels
  - Akin to bounding boxes for object detection
- ConvNet hyperparameters
  - 7 stacked layers (5 convolutional, 2 fully connected)
  - Fed thru ReLUs and smooth L1 loss regression layer
- Model spec
  - Keras on top of TensorFlow (Python)
  - Google GPU (8 vCPU Nvidia Tesla)

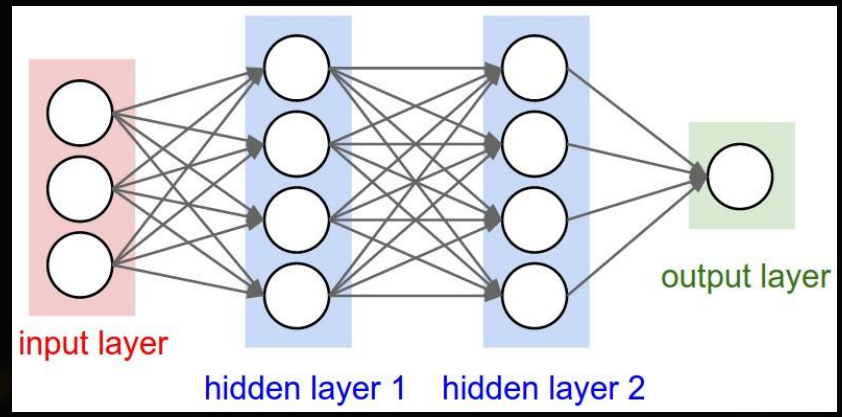

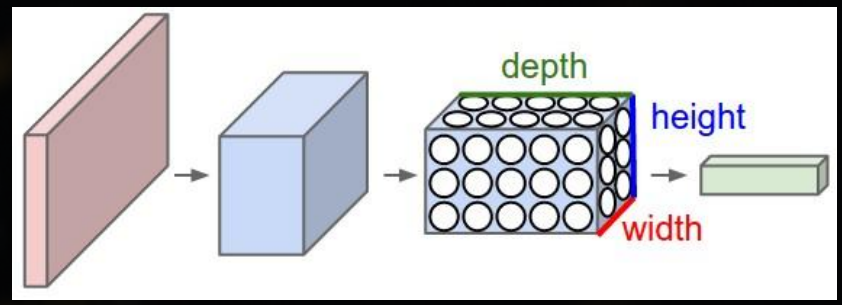

Illustration: Andrej Karpathy
CNNs: Szegedy, Toshev & Erhan, 2013

Prototype Evaluation

- More robust, less detectable transmission
- Learned locations correspond to locations that are more complex and “busier”
- Minimal Visual Dissimilarity
  - Distortion: peak signal-to-noise ratio
  - Capacity: byte Survivability
- Recovery rates worsen as hidden data size decreases

Data-Driven Red and Blue Teaming

A Picture is Worth A Thousand Words:

Deep Neural Networks for Social Stego

InfoSec ML Historically Prioritizes Defense

WILLIAM YERAZUNIS

Keeping the Good Stuff In: Confidential Information Firewalling with the CRM114 Spam Filter & Text Classifier

CLONEWISE - AUTOMATED PACKAGE CLONE DETECTION

Presented By: Silvio Cesare

DEFENDING NETWORKS WITH INCOMPLETE INFORMATION: A MACHINE LEARNING APPROACH

PRESENTED BY Alexandre Pinto

A SCALABLE, ENSEMBLE APPROACH FOR BUILDING AND VISUALIZING DEEP CODE-SHARING NETWORKS OVER MILLIONS OF MALICIOUS BINARIES

PRESENTED BY Joshua Saxe

FROM FALSE POSITIVES TO ACTIONABLE ANALYSIS: BEHAVIORAL INTRUSION DETECTION MACHINE LEARNING AND THE SOC

PRESENTED BY Joseph Zadeh

AN AI APPROACH TO MALWARE SIMILARITY ANALYSIS: MAPPING THE MALWARE GENOME WITH A DEEP NEURAL NETWORK

Konstantin Berlin | Senior Research Engineer, Invincea Labs, LLC

BOT VS. BOT FOR EVADING MACHINE LEARNING MALWARE DETECTION

PRESENTED BY Hyrum Anderson

TIME

Data-Driven Social Engineering

- Black Hat/DEF CON 2016
- Why Twitter?
  - Bot-friendly API
  - Colloquial syntax
  - Shortened URLs
  - Abundant personal data
- Machine grammar suffices
- 10k+ DoD accts targeted

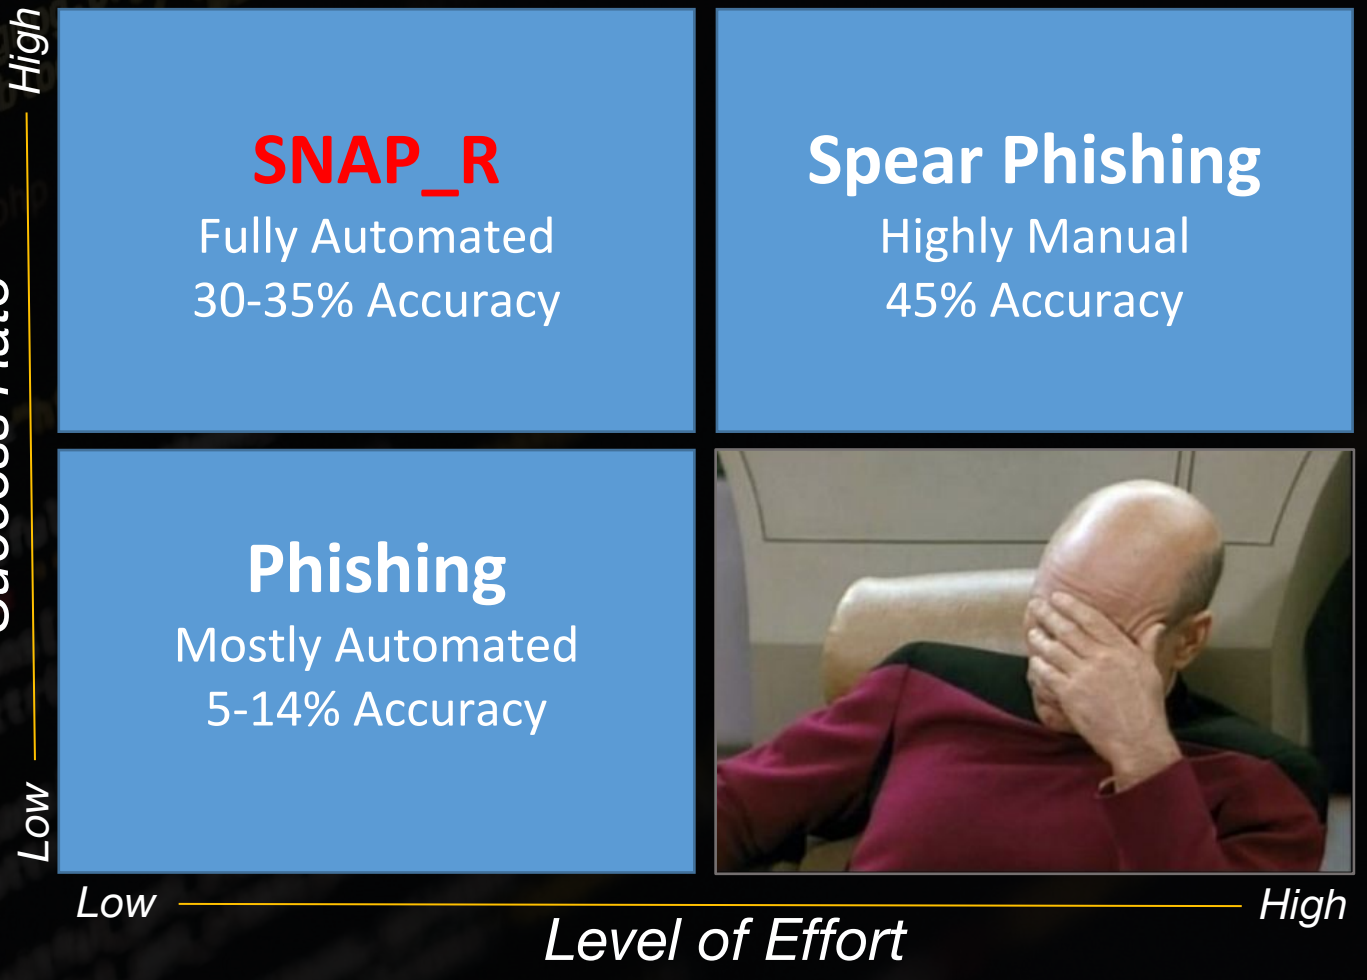

Red Team ML Rising

- Growing number of examples:
  - Micro-targeted social engineering
  - Password cracking
  - Captcha subversion
  - AV evasion
  - Steganography
- Offensive ML easier than defensive ML!
  - “Labeling Bottleneck” - unsupervised
- Success matters more for blue than red team
- Retreating barriers to entry
  - More open-source initiatives
  - Cheapening access to powerful machines (eg. GPUs)

Not to worry, though...

- Offensive ML is a positive development
- It will “keep us honest”
- Emerging defenses keep pace:
  - Semi-supervised learning
  - Adversarial learning
  - Transfer learning
  - Self-supervised reinforcement learning
- Ultimately improve security
- Faster this is realized, the better

Wrap Up

A Picture is Worth A Thousand Words:

Deep Neural Networks for Social Stego

Next Steps

- More social networks
- More stego (frequency domain)
- Video files (MP4, MOV, etc.)
  - Soon-to-be most popular
  - News Feed promoted
- Audio files (MP3)
  - Create custom MP3s w/ GarageBand
  - MP3s embedded JPEG insertion
  - ID3 Headers DC 24 SkyTalks Hosmer/Raggo www.python-forensics.org

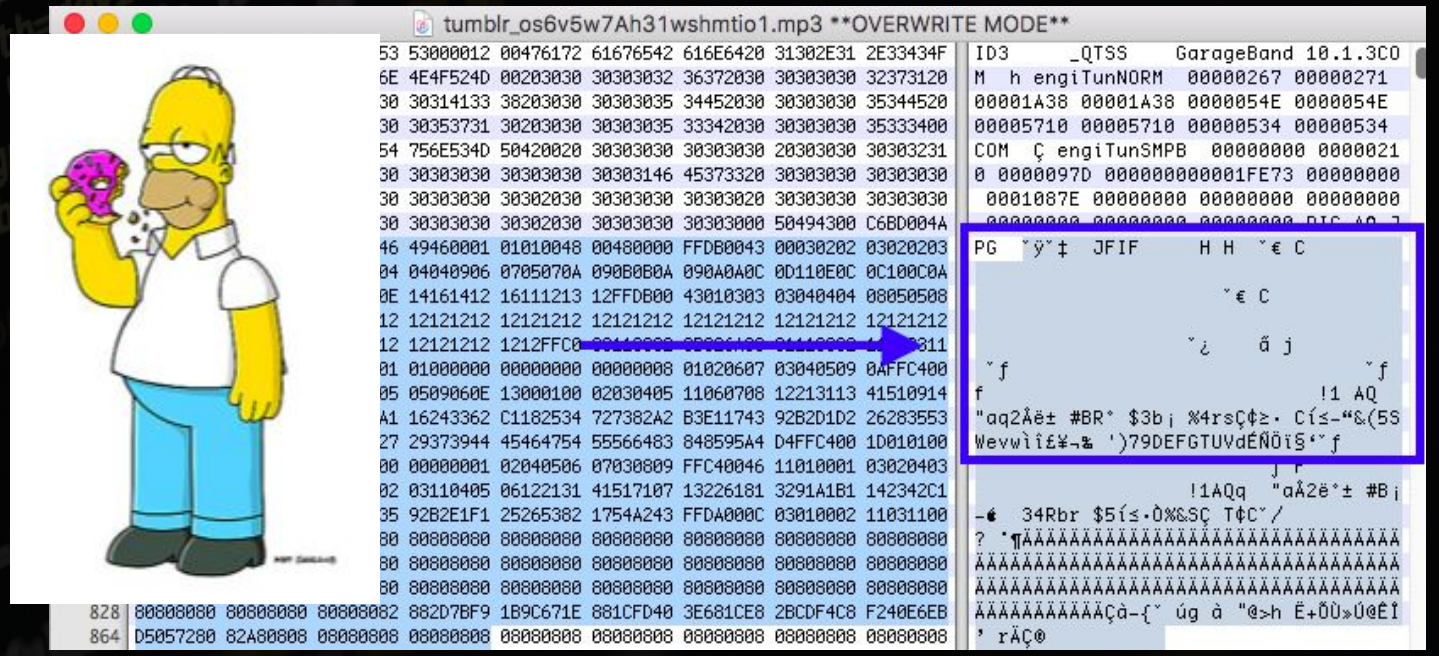

Mitigations

- More sophisticated and dynamic jamming techniques
- Anomaly/Histogram analysis - increased quantization
- Impermanence: delete by default
  - Ephemeral images a la Snapchat
- But generally, steganalysis is hard!
  - Variance in social networks add exponential complexity to identifying existence of stego and recovery of evidence - “know thy enemy”

Summary and Questions?

Philip Tully Mike Raggo

@phtully @datahiding

- Social networks and image hosting services can be orthogonally used to transmit data covertly
- Steganography can be automated despite distorting image upload side effects
- Offensive AI is cheaper and easier to implement than defensive AI
- Code released soon, PRs welcome