

CS 798: Digital Forensics and Incident Response

Lecture 6 - Steganography and Watermarking


Diogo Barradas

Winter 2025

University of Waterloo

Hiding in plain sight

Can you spot the difference?



Do they carry the same information?

- There is more than meets the eye!
 - One of the images encodes a secret message!



Decode

Hidden message!

The University of Waterloo (UWaterloo, UM, or Waterloo) is a public research university with a main campus in Waterloo, Ontario, Canada. The main campus is on 404 hectares (998 acres) of land adjacent to "Uptown" Waterloo and Waterloo Park. The university also operates three satellite campuses and four affiliated university colleges.[10][11] The university offers academic programs administered by six faculties and thirteen faculty-based schools. Waterloo operates the largest post-secondary co-operative education program in the world, with over 28,000 undergraduate students enrolled in the university's co-op program.[12] Waterloo is a member of the U15, a group of research-intensive universities in Canada.[13]

The institution originates from the Waterloo College Associate Faculties, established on 4 April 1956; a semi-autonomous entity of Waterloo College, which was an affiliate of the University of Western Ontario.[14] This entity formally separated from Waterloo College and was incorporated as a university with the passage of the University of Waterloo Act by the Legislative Assembly of Ontario in 1959.[2] It was established to fill the need to train engineers and technicians for Canada's growing postwar economy. It grew substantially over the next decade, adding a faculty of arts in 1960, and the College of Optometry of Ontario (now the School of Optometry and Vision Science), which moved from Toronto in 1967.[2]

The university is a co-educational institution, with approximately 36,000 undergraduate and 6,200 postgraduate students enrolled there in 2020.[4] Alumni and former students of the university can be found across Canada and in over 150 countries; with a number of award winners, government officials, and business leaders having been associated with Waterloo.[11] Waterloo's varsity teams, known as the Waterloo Warriors, compete in the Ontario University Athletics conference of the U Sports.

1. Steganography
2. Steganalysis
3. Watermarking

Steganography

Definition of steganography

- **Steganography:** Art and science of communicating in a way that hides the existence of a message
 - From the Greek words *steganos* and *graphy*
- Steganography takes one piece of (**secret**) information and hides it within another (**carrier / cover**)



Cybercriminals are known to use steganography

- Steganography is used for the concealment of communications in multiple crimes, e.g., terrorism, botnet management, data exfiltration, etc.

NEWS

Criminals using steganography tricks to manage ZBOT attacks

Researchers at Trend Micro have discovered examples of botnets being managed by steganography techniques.

Hidden file upload

Hackers used data exfiltration based on video steganography

November 29, 2014 By [Pierluigi Paganini](#)

 28

 My Page  Like  131

Security experts have detected an attack against a major firm that used a data exfiltration technique based on the video steganography.

Hidden file download

Steganography and terrorism: Why ISIS relies on it so much

Written by: [Vicky Nanjappa](#) Published: Monday, March 2, 2015, 11:17 [IST]

Hidden bidirectional communication

Cryptography vs. Steganography




Cryptography vs. Steganography



- **Cryptography:** is used to protect the contents of messages
- **Steganography:** is used to conceal the existence of messages

Steganography system model



- Wendy can be seen as a **warden**, and can be:
 - **Passive:** attempts to detect whether Y carries secret content
 - **Active:** modifies stego image Y into Y' in hopes of destroying the secret content

Desirable properties of steganography

- **Imperceptibility**
 - The resulting stego image appears innocuous and without perceivable visual artifacts
- **Security**
 - Must be able to resist active as well as passive attacks
- **Capacity**
 - Hiding capacity should be as large as possible

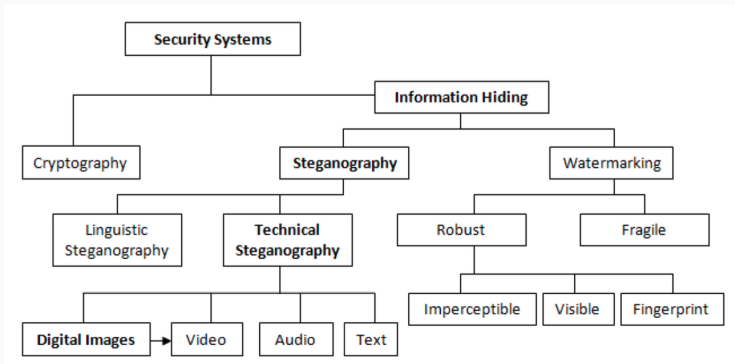
Steganography is not a new idea!

- Ancient Romans used to write using **invisible ink**
 - Based on various natural substances like fruit juices or milk
 - Messages would only appear when exposed to heat
- Still popular today! Think UV ink



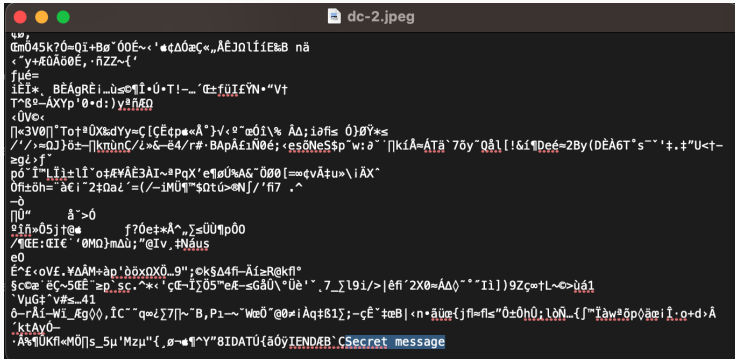
Digital steganography

- **Digital Steganography:** works by encoding secret bits in files, such as photos or audio files, with secret data
 - The secret message and the carrier message are digital objects



Basic example: exploit image file formats

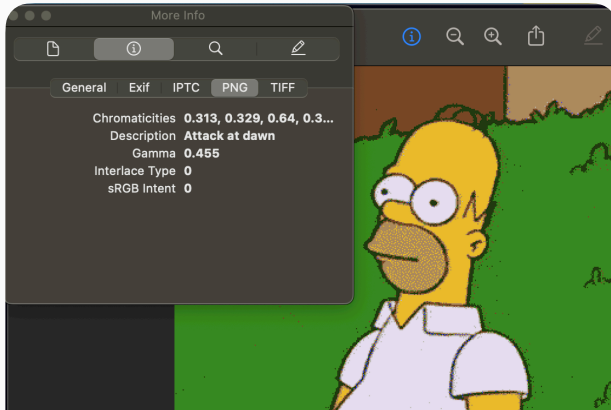
- Hide the secret into unused portions of the carrier file
 - echo "Secret message" >> dc-2.jpeg
 - Appends "Secret message" to the JPEG file
- A photo editor ignores anything coming after the image fields
 - However, the message is revealed once opened in a text editor



```
U00,
Gm045k70=Qi+B0^00E~<'!qΔ0æÇ«„ÅÊJQlÍíE&B nă
<'y+ÆUÂö0E,·ñZZ~{ '
fµé=
iÊI*, BEÁgREi...ú<0qI·Ú·T!-...´æ±fÜIEŸN·“V†
T^82-ÁXyp'0:d:)y.ñÆQ
<ÜV0<
[]«3V0[]"Tot±ÜX&dYy=Ç[ÇÊqþ«Å"}√<9~æ0î\% ÅΔ;iðfi≤ 0)ØŸ*≤
/'>/=QJ]ô±~[]kπünC/¿»&-ê4/r#·BApÂf1N0é;<es0Ne$Sp'w:ð~'[]kíÂ=ÂTä`76y`Qâ\!|!&í¶Deé=2By(DÊÂ6T*s~'†.‡"U<t-
z9¿>f~
pó`i~L_i±iÎ~o±Æ¥ÂÊ3ÂI~PqX`e[]ØÚ%A&~ÖØ0[=æqvÂ±u»\iÄX`
0fi±ðh="â€i`2±Qa¿'=(/-iMÜ¶""$Qtú>@Nj/'fi7 .^
~ð
[]Ü"      ä~>0
?iñ»05jt@#      f?0e‡*Å^„Σ≤ÜÜ[]p00
/¶EE:€IE'`0MQ}mΔú;"@Iv,‡Nâus
e0
Ê~f<0V£.¥ΔÂM+âp'ððxNXÖ...9";0k$Δ4fi-Âi≥R@kfi°
$ç0æ`êç~50Ê~≥p.sç.~*~*~'çE-iΣ05~0E-≤GâÜ\°Üè'`.7_Σ19i/>|éfi`2X0=ÁΔð~`"Ii}}9Zçω+L~ω>ùá1
`VµG±`v#≤..41
0-rÂi~Wi_ßg00,ÏC~`q=¿Σ7[]~"B,P1~`W00`@0#iÂq±B1Σ;-çÊ~±æB|<n·âÜæ{f fi=fi≤"0±0hÜ;1ðÑ...{j~i`âw#ôp0âæiî:q+d·Â
`ktAy0-
`Å%¶UKfi«M0[]s_5µ'Mzµ""{,0~!¶^Y"8IDATÜ{ä0ŸIENDÆB`CSecret message
```


Another example: append secrets to metadata headers

- Extended File Information (EXIF)
 - Standard used by digital camera manufacturers to store information in the image file, such as the make and model of a camera, the time the picture was taken, etc.
- PNG's description



Another example: append secrets to file extensions


- GIF's comment extension



G	I	F	8	9	a	4	L	0	0	0	0	f	3	"	>	>	f	U	"	0	0	^	U	D	Y	w	f	D	"	0				
0	f	D	I	w	U	w	D	3	i	^	w	U	3	"	>	>	w	U	0	f	U	I	w	f	0	0	f	D	3	i	0	w	0	
U	D	^	D	3	3	0	0	i	^	f	D	3	"	Y	0	w	I	^	f	f	3	3	a	f	D	>	w	f	U	3	3	a	f	
U	3	"	0	0	0	0	^	D	D	Y	^	f	w	D	D	"	"	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
0	!	p	g	S	t	e	g	a	n	o	g	r	a	p	h	y	a	m	o	n	g	o	t	h	e	r	r	a	r	e				
	d	i	s	c	i	p	l	i	n	e	s	i	s	h	o	n	o	r	e	d	t	o	b	e	d	e	s	c						
	r	i	b	e	d	a	s	b	o	t	h	a	n	a	r	t	a	n	d	S	c	i	e	n	c	e	f							
	i	e	l	d	.	0	!	ù	0	0	2	0	0	,	0	0	0	0	4	0	L	0	0	0	p	@	ε	p	H	,	0	0	h	


Exploiting the spatial domain

- A steganographer can insert the secret data in the spatial domain of the cover medium
 - For images, involves encoding information in pixels



Colour images

- 24-bit RGB image files
 - Each pixel encoded by 3 byte values for red, green, and blue



Grayscale images

- Pixels encode intensities in the form of shades of gray:
 - From black to white (ascending)
- Encoding defines how many bits per pixel: 8, 16, 32...
 - More bits means more levels of shade



A common digital steganography technique: LSB

- Least Significant Bit (LSB)
 - The LSB of a pixel is used to encode hidden information
- Example: Suppose we want to encode the letter A in the following 8 bytes of a 8-bit grayscale carrier image
 - “A” is 65 in ASCII 01000001 in binary

01011101	11010000	00011100	10101100
11100111	10000111	01101011	11100011

becomes

0101110 <u>0</u>	1101000 <u>1</u>	0001110 <u>0</u>	1010110 <u>0</u>
1110011 <u>0</u>	1000011 <u>0</u>	0110101 <u>0</u>	1110001 <u>1</u>

This can also be used in some audio formats (e.g., WAV)

What happens to the image?



Original



Stego image ('A' embedded)

- This works because:
 - Digital image or sound files can be altered to a certain extent without losing their functionality
 - Humans are unable to distinguish minor changes in image colors

Leveraging the RGB channels in color images

- Consider a 24 bit picture
 - Data to be inserted: character 'A': (01000001)
 - Host pixels: 3 pixel used to store one character of 8-bits


R	G	B
00100111	11101001	11001000
00100111	11001000	11101001
11001000	00100111	11101001

Embedding 'A'

0010011 <u>0</u>	1110100 <u>1</u>	1100100 <u>0</u>
0010011 <u>0</u>	1100100 <u>0</u>	1110100 <u>0</u>
1100100 <u>0</u>	0010011 <u>1</u>	11101001

Apply to LSB to color image carrier

- Tweaking the LSB is only a small change in image color
 - Pixel = 00100111, 11101001, 11001000
 - Pixel' = 00100110, 11101001, 11001000



Original



LSB modified to hide message

What happens if we encode on other bits?

- Different results in terms of capacity and added noise
 - More bits means higher capacity, but higher noise
 - Image distortions are more noticeable when we embed data in more significant bits



Original



4th LSB modified



Original



7th LSB modified



Original



6th LSB modified



Original



MSB modified

What kind of data can be used as payload?

- Pretty much anything!
 - An arbitrary sequence of binary data
 - Encrypted data too



LSB: The good, the bad, and the ugly

- **The good**

- Simple to implement
- Allows for large payload

- **The bad**

- Easy to figure out message if attacker knows the msg is there
- Vulnerable to statistical analysis

- **The ugly**

- Integrity is extremely frail
- Easy for attacker to corrupt the message
 - e.g., just randomize LSBs
- Vulnerable to unintentional corruption
 - e.g., image cropping, conversion to JPEG and back, etc.

Adding a key

- Detecting naively LSB-encoded messages is pretty easy
 - How do we make detection harder?
- **Shifting:** Simple approach that indicates an index $\neq 0$ of the first stego pixel to embed bit secrets
 - Example: `index_first_stego_pixel = key % 27`
 - The modulo operator ensures the shift is bound
- **Randomizing:** Use a pseudorandom number generator to spread the secret message over the cover image

Other digital steganography techniques

- **Substitute**

- Substitute redundant parts of a cover with a secret message
- Bit plane methods (LSB), palette-based methods

- **Transform**


- Embed secret info in a transform space of a signal, such as audio frequency domain

- **Distort**

- Store information by signal distortion and measure the deviation from the original cover in the decoding step

Substitute: Embedding in specific LSB planes

- BW image embedded into the image's 6th bit plane



Original Image





Image with another hidden image



Hidden image (in B6 plane)



Substitute: Embedding data with color palette tricks

- A **palette** is the set of available colors from which an image can be made
- Reducing the palette size may force secret messages to show up



Transform: Embedding data within audio frequency spectrum

- Principle: Encode digital files as sound waveforms (think a bit about ZX Spectrum audio cassettes)
- Other interesting example:
 - Encode images into (weird) sound files whose spectrograms look like these input images



Steganography tools

- Steghide (WAV, BMP, JPG...)
- S-Tools (GIF, JPEG...)
- Invisible Secrets (JPG, PNG, WAV...)
- Coagula
- Hiderman
- So many (even personal) others...

Steganalysis

What is steganalysis?

- The art and science of detecting hidden data
 - By identifying the existence of a hidden message, perhaps we can identify the tools used to hide it
 - If we identify the tool, perhaps we can use that tool to extract the original message
- Why is it important?
 - Prevent terrorist attacks
 - Catch people engaging in illegal activities



Types of steganalysis

- **Targeted steganalysis**

- Relies on knowing the method used to hide the data & using known distinguishing statistics to detect stego images
- May involve the reverse engineering of steganographic algorithms

- **Blind steganalysis**



- Most beneficial to forensics because it is not based on knowing the algorithm
- Most difficult because the type of images and method of hiding data are enormous and continuously changing
- The current trend is to develop a neural network using training images and multiple statistical features

Technical approaches: Structural detection

- **Inspection of file properties/contents**
 - E.g., content modifications, size difference, checksum / hash
- **Analysis of modification of file contents**
 - If you have a copy of the original file, it can be compared to the modified suspect/stego file
 - Many tools can help view and compare the contents of hidden file to identify inconsistencies and patterns, e.g., hex editor
- **Finding a signature of the steganography algorithm**
 - Reviewing multiple files may identify a signature pattern related to the steganography tool
 - Signature reveals program used to hide the message!
 - No original file necessary to compare it to



Example: Hiderman

- Consider a slightly sophisticated stego program - Hiderman
- After hiding a message with Hiderman, review the file with your favorite hex editor
 - Viewing the Header information (beginning of the file) we see that it is a Bitmap as indicated by the “BM” file signature




Case study: Hiderman

- Compare the original file to the stego file
 - Note the data appended to the file




Case study: Hiderman

- In addition, note last three characters "CDN": 43 44 4E in HEX




- Hiding different messages in different files with different passwords, ("CDN") is always appended: Signature found!




Technical approaches: Statistical detection

- Example of histogram analysis
 - Can be used to possibly identify a file with a hidden message
 - Histogram on the right has a very noticeable repetitive trend



original





stego

Technical approaches: Visual detection

- Detecting steganography by “eyeballing” it

Grainy pattern suggests the use of steganography



original

Inconclusive



carrier



steganogram



Least significant bits: black for LSB=0,
white for LSB=1

Watermarking

Cryptography, Steganography, and Watermarking

- Cryptography is about protecting the content of messages
- Steganography is about concealing the existence of messages
- **Watermarking** is about establishing identity of information to prevent unauthorized use:
 - They are imperceptible (most of the times)
 - They are inseparable from the works they are embedded in
 - They remain embedded in the work even after transformation



Examples: watermarks in bills

- Detect bill counterfeiting



Examples: machine ID codes in laser printers

- Some printers print yellow tracking dots on their output
- By decoding the tracking dots, the ID can be recovered



Examples: visible watermarks in images



Examples: visible watermarks in images

- **Copyright protection**
 - Embed owner info to prevent 3rd parties from claiming copyright
 - Requires a strong robustness
- **Copy protection**
 - Embed a watermark to disallow unauthorized copying
 - For example, a compliant DVD burner will not copy data that carries a “do not copy” watermark
- **Content authentication**
 - Embed a watermark to detect modifications to the cover
 - Requires a weak robustness

Steganography vs. Watermarking: Requirements

Steganography


- Robustness is not typically an issue
- Desired capacity is large
- Always invisible
- Dependent on file format

Watermarking

- Robustness is a big concern
- Can be visible or invisible
- Watermarks can be considered an extended data attribute

The magic triangle


- Trade-off between capacity, security, and robustness




Digital watermark

- A digital signal or pattern inserted into a digital image

Watermark Transmission:



Watermark Detection:




Simple recovery technique: Checksum embedding


- Recover the watermark by applying a checksum function to each pixel of auth image and check LSBs



Original Image with watermark embedded



Watermark




Altered Image



Reconstructed Watermark

Attacks to digital watermarks

- **Removal:** Remove the watermark, restoring the original
 - e.g., via content compression
- **Geometrical:** Unsync the watermark detector
 - e.g., via content rotation
- **Cryptographic:** Crack the watermark security methods
 - e.g., bruteforce key search
- **Protocol:** Attack the algorithms of a watermarking app
 - e.g., watermark inversion, copy attack



Takeaways

- Digital steganography is an increasingly used technique for concealing communications within criminal activities and is difficult to mitigate by investigators
- On the other hand, digital watermarking helps investigators to trace the real identity of digital media
- Both fields are relatively young, and research is ongoing in order to increase the security and robustness of these techniques

- **Textbook:**
 - Johnson – Chapters 1, 2.2-2.3, 3.1-3.2
- **Literature:**
 - Digital Image Steganography: Survey and Analysis of Current Methods. Signal Processing, Volume 90, Issue 3, March 2010
 - ML watermarking
- **Acknowledgements:**
 - Slides adapted from Nuno Santos's Forensics Cyber-Security course at Técnico Lisbon