

# Beyond Hamming Distance: Exploring Spatial Encoding in Perceptual Hashes

**Sean McKeown**

Edinburgh Napier  
UNIVERSITY



[/in/mckeowns87](https://www.linkedin.com/in/mckeowns87)



[s.mckeown@napier.ac.uk](mailto:s.mckeown@napier.ac.uk)

# Image Matching with Perceptual Hashes

*AKA Semantic Approximate Matching*

# Distances - Modified Image

**astronauts.png** vs **astronauts.jpg** ←

<u>Hash</u>	<u>Distance</u>
<b>ahash</b>	<b>0.0</b>
<b>whash</b>	<b>0.0</b>
<b>dhash</b>	<b>0.0</b>
<b>phash</b>	<b>0.0</b>
<b>blockhash</b>	<b>0.0</b>

**astronauts.png** vs **astronauts\_edit.png** ←

<u>Hash</u>	<u>Distance</u>
<b>ahash</b>	<b>0.015625</b>
<b>whash</b>	<b>0.0</b>
<b>dhash</b>	<b>0.046875</b>
<b>phash</b>	<b>0.093750</b>
<b>blockhash</b>	<b>0.023438</b>



**Image Diff**  
(compression)



**Image Diff**  
(edit)

# Distances – Different Images

a.jpg vs b.jpg

Hash	Distance
average_hash	0.703125
whash	0.718750
dhash	0.546875
dhash_vertical	0.609375
phash	0.625000
phash_simple	0.578125
blockhash	0.601563



# Calculating Distance

## Hamming Distance

- ◇ XOR bit strings, count 1s

## Normalised Hamming Distance

- ◇ Divide by hash length
  - ◇ Result is between 0 and 1
- ◇ Captures **global** difference between hashes
  - ◇ Positional information is completely lost
  - ◇ Often not important, or just not leveraged?

Hamming Distance = 8

A	1	0	0	1	0	1	1	0	0	1	0	0	1	0	0	1
B	0	1	0	1	0	0	1	0	1	1	1	0	0	1	0	0
XOR	1	1	0	0	0	1	0	0	1	0	1	0	1	1	0	1

Hamming Distance = 8

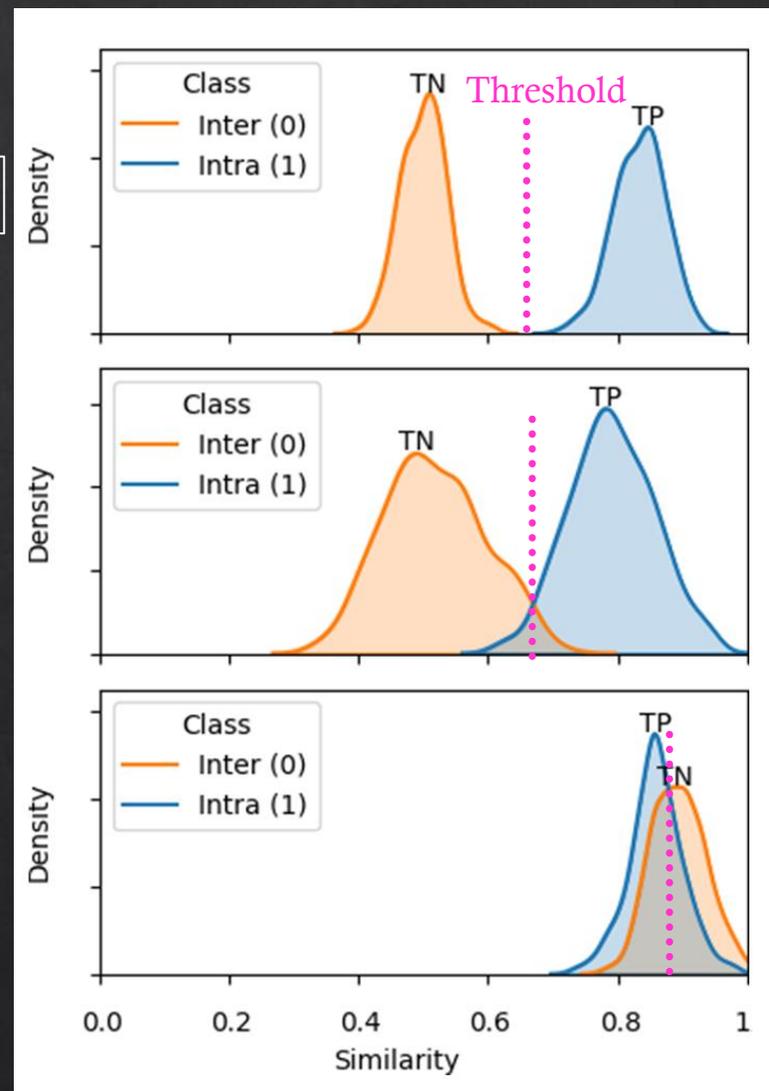
A	1	0	0	1	0	1	0	1	0	1	0	0	1	0	0	1
B	0	1	1	0	1	0	1	1	0	1	1	0	1	0	0	1
XOR	1	1	1	1	1	1	1	0	0	0	1	0	0	0	0	0

# Macro Goal: Separate Distance Distributions

## Unrelated Images vs. Modified Originals

- ◆ Set a **Distance Threshold** that lets us determine if an image pair are a:
  - ◆ **Match**
  - ◆ **No Match**
- ◆ Overlap causes **False Positives / False Negatives**
- ◆ Focus of prior work is typically the Perceptual Hash, not the Distance Metric

Best Case



Some Errors

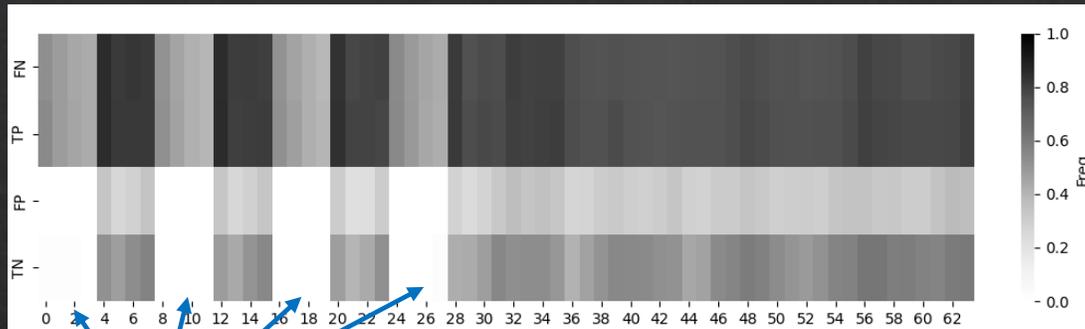
Terrible  
*Flip a  
Coin*

Does Positional Information Matter?

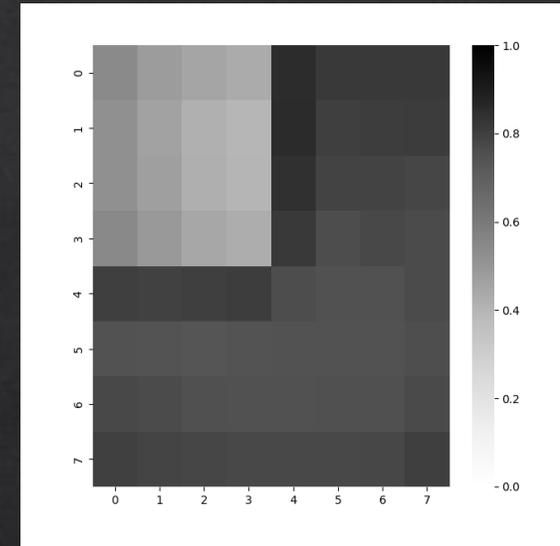
# Approach

- ◆ Determine if **redundant data** in an image is identifiable in output Perceptual Hashes
  1. Insert constant data at fixed positions (e.g., top-left, border, watermark)
  2. Compare original images to modified images -> aggregate over a dataset
  3. **Re-weight** hash bits based on their contribution towards correct classifications
    - ◆ **Low-weights** = spatially encoded redundant data (**no discriminatory power**)
- ◆ Similar approach for transposition (mirroring, rotation), and crop
- ◆ Tested on spatial-domain hashes (ahash, dhash) and DCT-based hashes (pHash, PDQ)

# It works: Top-left Redundant Data Insertion



Low-weights indicate these bits are not useful for discrimination in the classification process

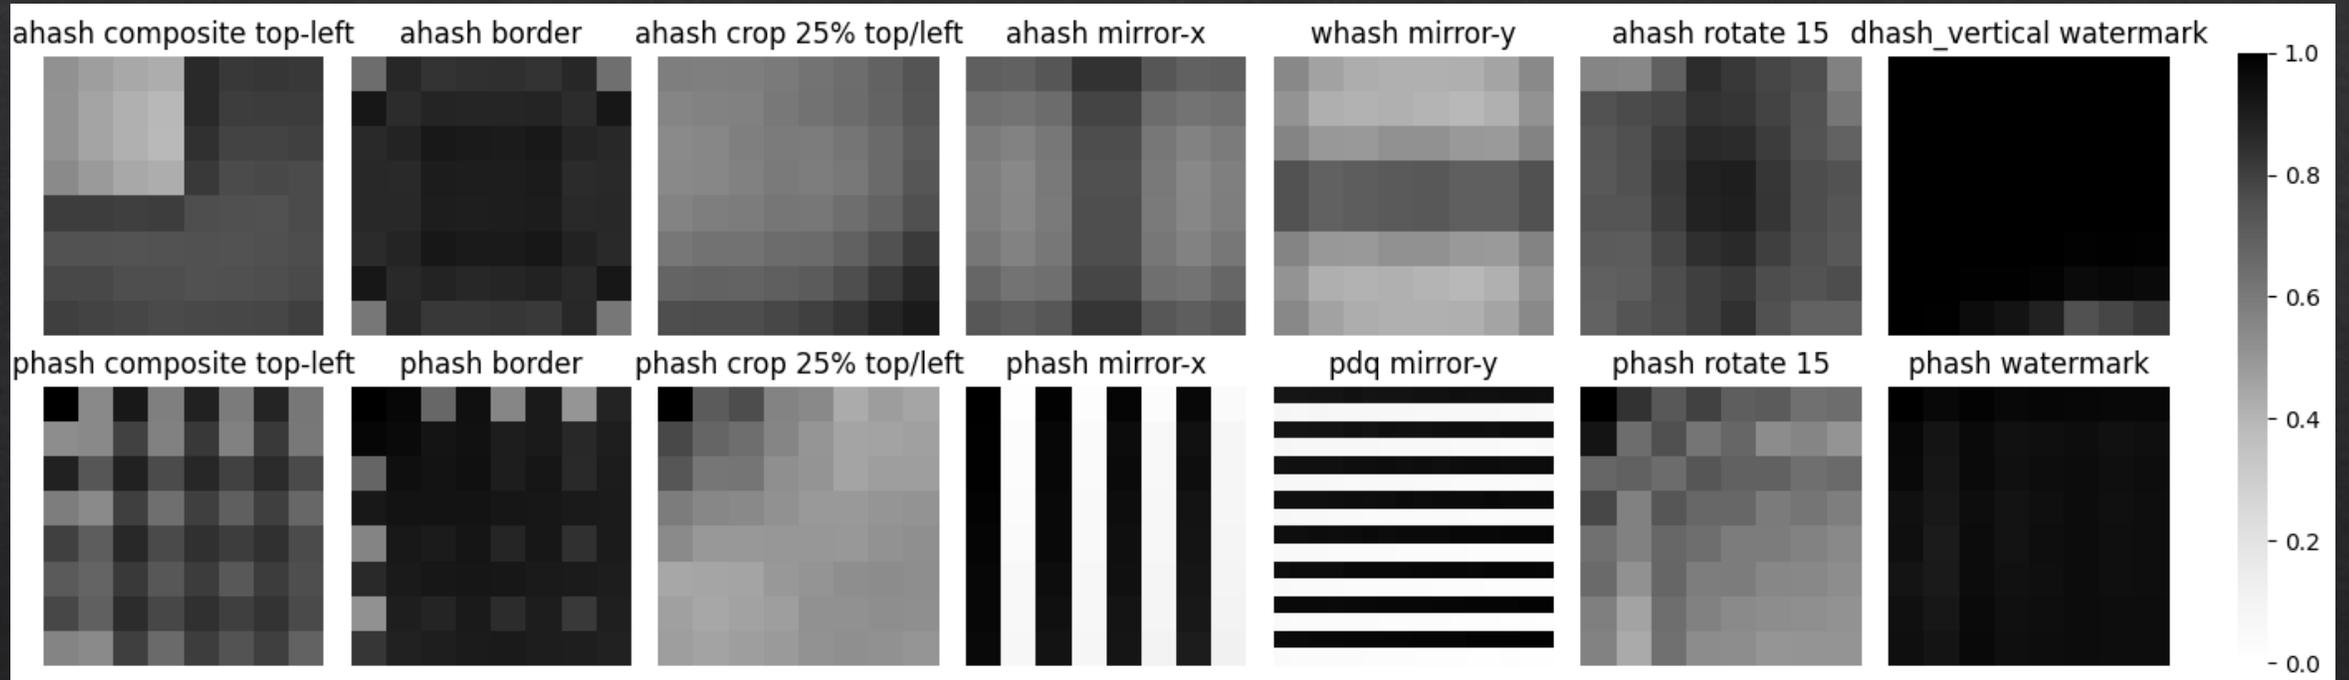


Wrapping to a square (TP only) makes the pattern clearer

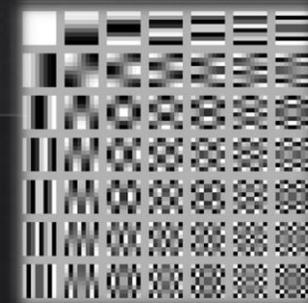
# Patterns for a Range of Changes

Spatial-domain

Weight



DCT (Frequency-domain)



# Leveraging Hash-bit Positional Encoding

# New Metrics!

- ◇ Common metrics align strongly with Normalised Hamming Distance
  - ◇ Scipy.Spatial.Distance (e.g., L1, Earthmover, Manhattan, Cosine)
- ◇ New metrics needed to capture hash “locality”

## Normalised Convolution Distance

- ◇ Convolutions on XOR of hash matrices

## 2D N-Grams

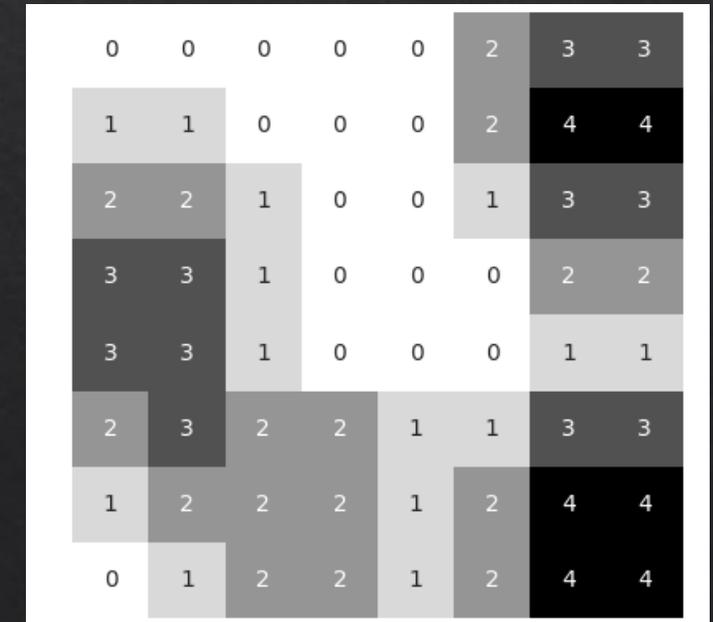
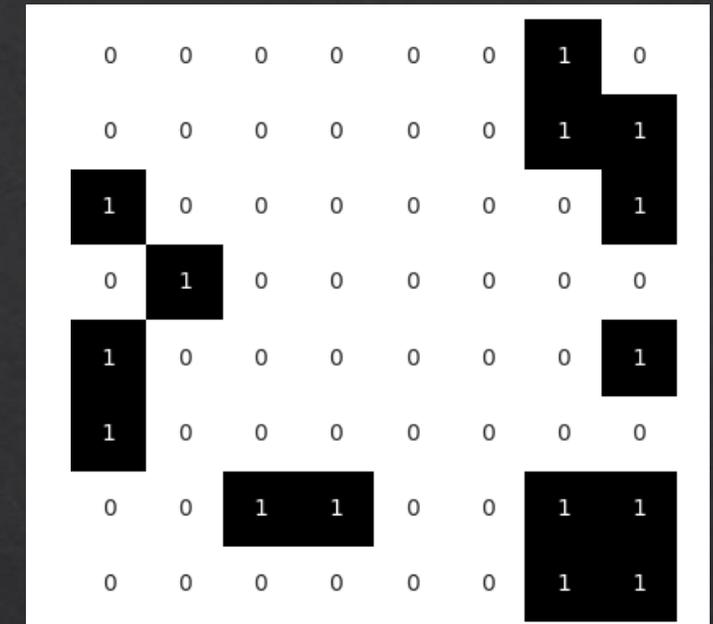
- ◇ Sliding windows to capture locality, compare windows between hashes

## Hatched Matrix

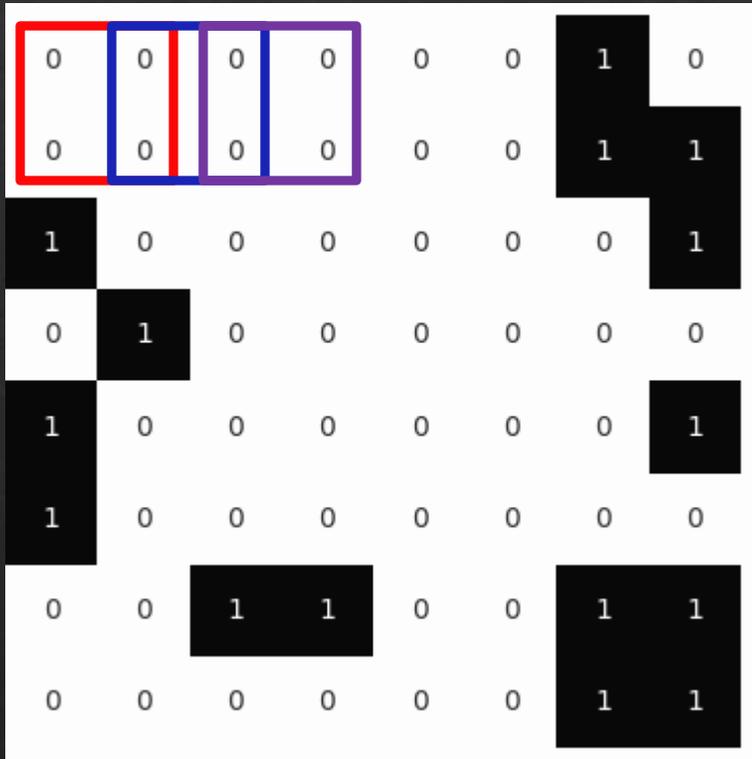
- ◇ Take advantage of the row/column pattern for pHash/PDQ

# Convolution Distance

1. **XOR** input hash matrices
2. **Convolve**
  - ◇ Multiple kernel configurations
  - ◇ Settled on (4,4) matrix of ones
  - ◇ Captures weighty clusters of change
3. **Sum** all matrix elements
4. **Normalise** by maximum possible distance
  1. Dependent on kernel, matrix size

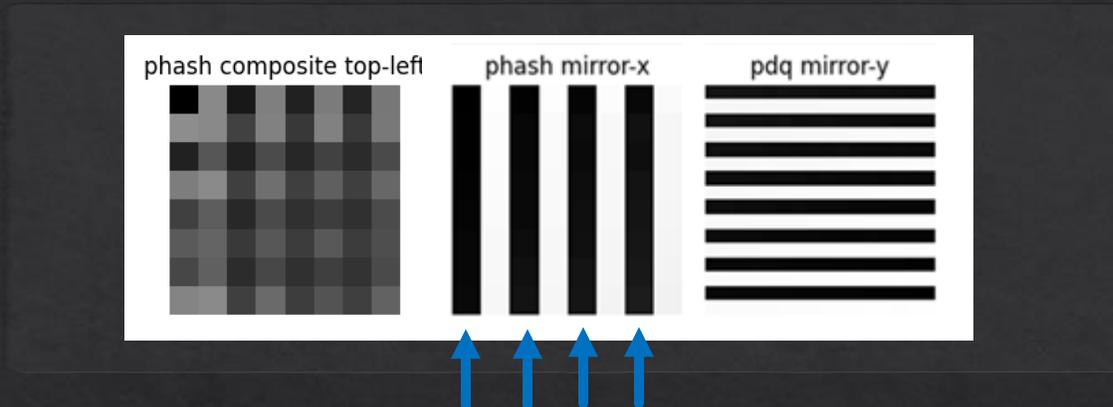


# 2D N-Grams



1. **Slide a two-dimensional window** over all elements **of each hash**
  - ◇ Overlapping windows
  - ◇ Various sizes tested, 2x2 chosen
2. **Flatten and concatenate** all windows (single array for each hash)
3. **Calculate Cosine Distance** between both arrays

# Hatched Matrix Distance



- ◇ DCT-based hashes largely generate Row/Column patterns of coefficient weights
  
- 1. **Extract rows/columns** for each hash into their own arrays
- 2. **Compare odd/even rows/columns** (Hamming distance)
- 3. Take **minimum** value for row/column, then **average**
  - ◇ Biases towards similar rows/columns
  - ◇ Assumes this isn't by accident!

# Evaluation

- ◇ Compare vs. Baseline: Normalised Hamming Distance
- ◇ 250k subset of Flickr 1 Million
- ◇ Compare modified image to original: %pt difference to Hamming AUC
- ◇ Validated inter-image (no-match) distributions, covered in paper

# Results

## Hatched Matrix

- ◇ **Mirrored** images go from a weakness to a strength for pHash and PDQ (+49 %pts). Small benefit on **rotation**
- ◇ Largely detrimental to spatial hashes

## Convolution Distance

- ◇ Reasonably large spatial hash benefit for **rotation** (+2.5 - 8.4 %pts)
- ◇ Some gains in top-left insertion (+0.4 – 4.7 %pts)

## 2D N-Gram

- ◇ Disappointing, trade-offs are as large as gains

Hash	Trans.	AUC	%pt Diff to Hamming		
		Hammm.	Conv4.4	Hatch	2gram
ahash	CropTL	0.829	0.2	0.0	-5.4
	MirrorX	0.766	-0.3	-2.6	-3.6
	Rotate	0.919	<b>2.5</b>	-0.1	-0.8
	Border	0.971	0.7	-0.3	0.6
	CompTL	0.527	<b>3.8</b>	0.5	-4.8
dhash	CropTL	0.641	-0.6	-0.7	-4.9
	Mirr.X	0.618	-1.1	-3.4	<b>1.0</b>
	Rotate	0.808	<b>7.6</b>	-1.0	<b>4.3</b>
	Border	0.995	0.4	-0.3	0.1
	CompTL	0.992	0.4	-0.4	-2.9
dhash vertical	CropTL	0.646	-0.4	-0.4	-4.3
	MirrorX	0.801	-0.1	-2.9	-6.5
	Rotate	0.780	<b>8.4</b>	-1.0	<b>3.6</b>
	Border	0.992	0.6	-0.4	0.2
	CompTL	0.989	0.4	-0.6	-5.5
Neural hash	CropTL	0.996	0.0	-0.1	-0.4
	MirrorX	0.930	-0.1	-0.4	-1.0
	Rotate	0.988	0.0	-0.2	-0.4
	Border	0.999	0.0	0.0	0.0
	CompTL	0.844	0.9	-0.7	<b>2.7</b>
PDQ	CropTL	0.527	-0.2	-0.1	-0.9
	MirrorX	0.515	0.9	<b>48.5</b>	<b>1.6</b>
	Rotate	0.502	<b>1.7</b>	<b>3.0</b>	<b>2.7</b>
	Border	1.000	0.0	0.0	0.0
	CompTL	1.000	0.0	0.0	0.0
phash	CropTL	0.586	-2.5	-0.2	-0.3
	MirrorX	0.496	<b>2.3</b>	<b>49.1</b>	<b>1.1</b>
	Rotate	0.675	-0.1	<b>1.5</b>	-0.9
	Border	1.000	0.0	0.0	0.0
	CompTL	0.944	<b>2.4</b>	0.6	<b>1.3</b>
whash	CropTL	0.821	0.3	0.0	-2.1
	MirrorX	0.744	-0.1	-2.9	-1.7
	Rotate	0.904	<b>3.1</b>	-0.2	<b>1.1</b>
	Border	0.921	<b>1.4</b>	-0.5	<b>3.6</b>
	CompTL	0.604	<b>4.7</b>	0.2	-3.4

# Future Work

- ◇ Explore alternatives based on observed patterns
  - ◇ Crop isn't helped much with the three proposed algorithms
- ◇ Frequency domain:
  - ◇ DCT low-frequency coefficient weighting (incorporate into Hatched Matrix or other)
- ◇ Spatial domain:
  - ◇ Centre of the image should take more importance

# Thank You

## Perceptual Hash Evaluation Framework:

<https://zenodo.org/records/10363151>

<https://github.com/AabyWan/PHASER>

## Prior Work:

S. McKeown and W. J. Buchanan, 'Hamming distributions of popular perceptual hashing techniques', *Forensic Science International: Digital Investigation*, vol. 44, p. 301509, 2023.

S. McKeown, P. Aaby, and A. Steyven, 'PHASER: Perceptual hashing algorithms evaluation and results-An open source forensic framework', *Forensic Science International: Digital Investigation*, vol. 48, p. 301680, 2024.



/in/mckeowns87



s.mckeown@napier.ac.uk