




noRECOGNITION

Today, over 100 cameras logged your beautiful face.

You didn't opt in. You can't opt out.

Facial recognition is everywhere.

And it is being used against us.

117M+

Your drivers license photo is probably
in a police face database

higher false match rate for
darker-skinned individuals

10-100x

6,000+

police departments on the Flock
network

Sold to catch car thieves.
Rebuilt into a surveillance state.

140,000 monthly police users. Every passing car, logged. No warrant.

FOOTAGE TAKEN FROM AN **ALPR** CAMERA



- **Face identification**
Matched to an identity database
- **Emotion detection**
Inferred from micro-expressions
- **Screen reading**
Phone content captured

This is now.

Video credit: Benn Jordan
<https://www.youtube.com/@BennJordan>

Trusted to serve. They served themselves.

Sedgwick, Orange City, Joplin. At least 18 documented cases.
So far...



Hackers have always
been the counterbalance.





Hey I'm bill, I am a **HAKCER**

Bill Swearingen — ex-CISO, ex-red-team lead, founder of SeckC.

CISSP · CISM · CMMC-RP · TS/SCI



SECKC

The World's largest monthly hacker meetup

IT
STARTED AS
A GREAT
IDEA





A camera is just a **parser**.

It parses pixels into objects. Nothing more.

- Every parser in the history of computing has been exploitable when an attacker controls the input.
- SQL injection works because the database parser cannot distinguish data from commands.
- **I control what appears on camera.**

I control the input, can I impact the output?

Status

Adversarial Pattern Fuzzer - v0.5 Refactored (Classy)

Overall Stats


Epoch: 6

Total Anomalies: 11730

Epoch Rate: 737.7 tests/min

Stop Fuzzer: Press Ctrl-C

Anomaly Preview



Epoch Progress

Epoch 6 14% • 0:50:35 < 8:17:45 • 12.3 tests/sec

Epoch 5: Rate: 717 tests/min, Face Anomalies: 10820, Person Anomalies: 27885

Epoch 4: Rate: 561 tests/min, Face Anomalies: 7858, Person Anomalies: 21078

Epoch 3: Rate: 406 tests/min, Face Anomalies: 4897, Person Anomalies: 14372

Epoch 2: Rate: 662 tests/min, Face Anomalies: 2153, Person Anomalies: 7692

Epoch 1: Rate: 410 tests/min, Face Anomalies: 270, Person Anomalies: 1811

Anomaly Log

✦ Added 'saliency_eye_attack_seed8334109' to priority list.

Anomaly: saliency_eye_attack_seed8334109 (PERSON_LOST (Found: 0, Baseline: 1))

✦ Added 'trypophobia+gradient_seed4464024' to priority list.

Anomaly: trypophobia+gradient_seed4464024 (PERSON_LOST (Found: 0, Baseline: 1))

✦ Added 'qr_code+tiled_logo_seed1671641' to priority list.

Anomaly: qr_code+tiled_logo_seed1671641 (PERSON_LOST (Found: 0, Baseline: 1))

I had
invented something
AMAZING.
My friends were
quick to help.



JOSHUA MARPET WAS RIGHT

I had no freaking clue what I was doing

Now I know what you are thinking: this guy is a genius.

Most of this ground was already walked. They walked so I could run, and I am grateful they did.

THE ARTISTS

Adam Harvey, CV Dazzle (2010)

Makeup and hair styling designed to defeat face detection.

Adam Harvey, HyperFace (2017)

Textile printed with decoy face patterns to saturate detectors.

Kate Rose, Adversarial Fashion (2019)

Garments printed to inject false plate reads into ALPR systems.

THE PRODUCTS

Cap_able, Manifesto (2023)

Knitwear with adversarial patterns; 60% evasion claimed against one detector.

Urban Privacy

Anti-surveillance clothing offered at roughly 35 euros.

Reflectacles

Infrared-blocking eyewear aimed at camera and depth sensors.

THE RESEARCH

Thys et al. (2019)

Printed patch defeats YOLOv2 person detection.

Xu et al. (2020)

Adversarial T-shirt: physical attack on person detectors under deformation.

Wu et al. (2020)

Invisibility cloak: wearable attack on object detectors.

THE MEASUREMENT GAP

Claims are typically measured against a single object detector, in sample, without a control. Independent review of the category scored 6/10; of thirty published methods, twelve were ever tested against a real system.

Detection Models

Three distinct capabilities, three different questions



PERSON

How many people are there?

Detects human bodies in the frame.
Produces a count, not an identity.



FACE

How many faces are there?

Locates and isolates faces in the frame.
Still produces no identity.



FACIAL RECOGNITION

Do we know this face?

Matches a detected face against a gallery
of known identities.

Confidence

Every detection is a score, not a fact

1 PERSON DETECTED

Flock - YOLOv5 320 0.82



PERSON

A false positive adds a bounding box. Thresholds are commonly set low, since the cost of a missed presence exceeds the cost of a spurious one.

FACE

A face below threshold is never passed downstream. Threshold choice here constrains every recognition result that follows.

FACIAL RECOGNITION

A false match asserts an identity. Thresholds must be high, because the cost of error is attribution to a specific person.

Target	What it actually does	Closest open
Axis P1465-LE, Q1656-LE	On-device object detection	YOLOv5s or SSD-MobileNetV2
Intel OpenVINO (legacy smart city)	Face detection	MTCNN (legacy), SSD face-detection-retail (current default)
Hikvision AcuSense	Human/vehicle detection, <i>not</i> face	SSD-MobileNetV2 or YOLO -class
Dahua WizSense, Hikvision MinMoe	Face recognition (1:N)	RetinaFace + ArcFace (InsightFace buffalo_L)
Avigilon AI NVR	Person detection + appearance re-ID	PeopleNet (DetectNet_v2 / ResNet34) + OSNet-class re-ID
Axon Body 4, Fleet 3	Person/object detection, cloud-side, <i>not</i> recognition	YOLOv8n -class detector
Axon Redaction Assistant	Face/head + plate detection, <i>not</i> recognition	RetinaFace or SSD face detector
Clearview AI, state surveillance	Face recognition (1:N)	ArcFace patent hints at FaceNet-style metric learning

A security camera is mounted on a black pole, positioned below a blue awning. The camera is black and has a lens with a circular pattern of small lights. The background is a clear blue sky.

There are cameras everywhere
But y'all only care about defeating flock.

Every persona now runs a gauntlet of 11 models.

Five person detectors, four face detectors, two recognition models. One pass became eleven.

PERSON

5 MODELS

How many people are there?

-
- P1 YOLOv8n
 - P2 YOLOv5s
 - P3 SSD-MobileNet
 - P4 ResNet34-SSD
 - P5 YOLOv5 (added July 2026)

FACE

4 MODELS

How many faces are there?

-
- F1 InsightFace
 - F2 Facenet-MTCNN
 - F3 MTCNN
 - F4 RetinaFace

FACIAL RECOGNITION

2 MODELS

Is this face in the database?

-
- R1 ArcFace
 - R2 Facenet

THE COST

Each persona is scored against all eleven models rather than one. Every sweep now carries eleven times the inference load, and research throughput falls in proportion.

I'M GOING TO NEED MORE POWER



31.7 Million

PATTERNS TESTED

534.6K

ANOMALIES
DISCOVERED

1.7% of patterns tested

480.7k

MULTI-MODAL
ANOMALIES

89.9% of anomalies discovered

85

EXTREME
ANOMALIES

0.016% of anomalies discovered

31.7M patterns → 534.6K anomalies → 480.7K multi-modal → 85 extreme

Each pattern scored against the full gauntlet of 10 models.
The highest performers pushed into an evolution engine,
picking the strongest characteristics and breeding them.

**1000s of HIGHLY EFFECTIVE
PATTERNS FOUND**

I AM A GENIUS




192GB

Then two Blackwells showed up.

It's time to change tactics





https://www.youtube.com/watch?v=L_4BPjLBF4E

31.7 Million

PATTERNS TESTED

534.6K

**ANOMALIES
DISCOVERED**

1.7% of patterns tested

480.7k

**MULTI-MODAL
ANOMALIES**

89.9% of anomalies discovered

85

**EXTREME
ANOMALIES**

0.016% of anomalies discovered

31.7M patterns → 534.6K anomalies → 480.7K multi-modal → 85 extreme



All of it became training data.

Every win and every loss now trains a deep reinforcement learning system.

588

PERSONAS

Nearly twenty times the
corpus: broader coverage of
body types, ages, and skin
tones.



Barry



Carrie



Gary



Larry



Mary



Sherry





One image. 6,300 guesses. Thirteen numbers each.

What Flock's own model file says: flock_yoloV5.tflite, output shape [1, 6300, 13]

1

one image per pass

The camera runs a single frame at a time.

6,300

candidate boxes

Grids 40x40 + 20x20 + 10x10, three anchors each.

13

numbers per candidate

4 box + 1 objectness + 8 class scores.

The thirteen numbers, and the one that decides whether you exist

x	y	w	h	obj	c1	c2	c3	c4	c5	c6	c7	c8
4 box coordinates				OBJECTNESS	8 class scores: person, car, plate...							

Objectness answers one question: is there a thing here at all? Score = objectness x P(person); a box is drawn only if score ≥ 0.75 .

YOLOv5 — 320, what Flock ships

flock_yoloV5.tflite

in 320 x 320 x 3 → [1, 6300, 13]

One scalar gates every detection. LOCAL → the soft target.

SSD — 300, the other build

flock_small_tf.tflite · .dlc (SNPE)

in 300 x 300 x 3 → [1, 1917, 11]

10 classes + background, no objectness. GLOBAL → the wall.

The classes are just output channels: **person** · car · truck · bus · trailer · motorcycle · bicycle · **licensePlate**

The plate reader was never bolted on: licensePlate is a class, with its own post-filter areaThreshold = 0.00040690104

What the adversarial pattern is actually doing

YOLOv5-320, the objectness family: one scalar per candidate decides whether a box is drawn

01

What the filters expect

The backbone is an edge and texture detector, trained on natural photographs. Its filters are tuned for skin, cloth folds, hair and shadow.

02

What the pattern is

Dense, periodic, high-frequency structure. No filter in the backbone is matched to it, so the features that fire carry no evidence of a body.

03

What happens to the scalar


In the cells the garment covers, objectness falls. That is the single number answering: is there a thing here at all?

04

What the pipeline then does

score = objectness x P(person) collapses. Under the 0.75 keep-line no box is drawn, so NMS, tracking and upload never run.

MEASURED, HELD-OUT IDENTITY



Same body. Same camera. Same model. Only the texture changed.

Occlusion-subtracted against a solid black panel of the same geometry, so this is the pattern, not coverage.

Objectness is one scalar over a region a garment can physically reach.

WHY SSD WALLS

1917 anchors each run class against background. The garment reaches the torso boxes; whole-body and background boxes still see you, and one survivor keeps you detected.

CLEAN VERSUS PATTERN, WITH A SOLID-BLACK CONTROL

THE 0DAY SLIDE

P1 (YOLOv8n)
garment: v2_LARGE_JACKET persona: p2_00010

DEFEATED



CLEAN (detected)
clean person conf: 1 (0.92) -> with pattern: 0 (0.00) | black control: 1 (0.94)

WITH PATTERN

P2 (YOLOv5s)
garment: v2_LARGE_JACKET persona: p2_00012

DEFEATED



CLEAN (detected)
clean person conf: 1 (0.87) -> with pattern: 0 (0.00) | black control: 1 (0.81)

WITH PATTERN

P3 (SSD-MobileNet)
garment: v16_LARGE_FULLBODY persona: p2_00021

DEFEATED



CLEAN (detected)
clean person conf: 1 (0.98) -> with pattern: 0 (0.00) | black control: black 1 (0.97)

WITH PATTERN

P4 (ResNet34-SSD)
garment: v4_MEDIUM_HOOD_UP (head/torso) persona: p2_0000

DEFEATED



CLEAN (detected)
clean person conf: 5 (0.69) -> with pattern: 0 (0.00) | black control: bg-fill 0.81

WITH PATTERN

FLK (Flock Falcon v2)
garment: v1_LARGE_HOODIE persona: p2_00002

DEFEATED



CLEAN (detected)
clean person conf: 1 (0.89) -> with pattern: 0 (0.00) | black control: 1 (0.85)

WITH PATTERN

F1 (SCRFD)
garment: v6_MINIMAL_NECK_GAITER persona: p2_00008

DEFEATED



CLEAN (detected)
clean faces: 0.801 -> with pattern: 0.000 | black control: 0.781
CERTIFIED held-out 8/10, occ-sub +0.587 vs same-shape SOLID BLACK. Re-confirmed this run: 8/10 held-out personas_v2 gaiter, mean occ-sub +0.56

WITH PATTERN

F2 (MTCNN)
garment: v6_MINIMAL_NECK_GAITER persona: p2_00018

DEFEATED



CLEAN (detected)
clean faces: 1 (1.00) -> with pattern: 0 (0.00) | black control: 1 (1.00)

WITH PATTERN

F4 (RetinaFace-R50)
garment: patterned hoodie, hood up persona: p2_00026

NO EFFECT (occ-sub -0.000)



CLEAN (detected)
clean F4 conf: 0.999 -> with pattern: 0.999 | same-footprint BLACK hoodie control: 0.999
MEASURED in-session on F4_retinaface_clean onnx: sanity gate passes (solid black over the face -> 0.035).
Hoodie footprint = 10-75% frame coverage. This is NOT the certified halo footprint.
The F4 file on disk is the COWB file (evade_face_cowl_universal.p2). The halo file was never saved.
Certified sub-0.02 19/10 backdoor belongs to the -1.5% nonacc-loading RMLO ring and is NOT reproduced by this asset.

WITH PATTERN (still detected)

THE FINDING THAT CHANGED EVERYTHING

Certain patterns changed more than the confidence.

Some caused the bounding box to **move**.

Others caused unexpected configuration changes **TO THE CAMERA**.

WHERE THE PROOF STANDS

Digital, and held to a stricter bar than the field.

THE TESTING CRITERIA

■ Sealed sweeps

Train, select and report kept separate, so nothing is scored on the data that produced it.

■ Held out, multi-angle, 500+ people

More than 500 distinct personas, unseen at training time, across multiple view buckets and an expectation over transformations.

■ Occlusion subtracted

The pattern must beat a solid black panel of the same geometry, or the result is discarded.

LIVE FABRIC TESTING

Just begun.

Printed fabric trials are under way. Nothing physical is claimed yet: the digital results above do not transfer until fabric passes the same occlusion-subtracted bar.

Don't like this? Blame nvidia

DEMO TIME

MODEL UNDER TEST: Flock - YOLOv5 320

1 PERSON DETECTED

Flock - YOLOv5 320 0.80 SEEN x1

Flock - YOLOv5 320 0.80



THE NEXT QUESTION

A person already disappears from the fielded Flock detector.

Can I make a car disappear?

IN PARTNERSHIP WITH

DONUT MEDIA

Vehicles, camera time and a test environment for the plate and vehicle recognition work.

Same bar: controls in loop, occlusion subtracted.

DONUT LIVERY on 2009 Toyota Yaris

Flock camera detects a car at 0.25 or above. This wrap keeps it at 0.02-0.05.



OK GREAT WHERE DO I BUY ONE

I am opening a **Kickstarter** to fund the build-out.

- **Test fabrics**

Substrates, weaves and inks, to find what actually holds a pattern.

- **Dye-sublimation printer**

In-house printing so an iteration takes hours instead of weeks.

- **More cameras**


Fielded hardware beyond the bench, so results are measured through real lenses.

Same standard: controls in loop, occlusion subtracted, results published.



Scan to back the build

Kickstarter, launching with this talk

A woman wearing a green baseball cap and glasses, holding a brown chicken. The image is semi-transparent and serves as a background for the text.

noRECOGNITION

Visible to humans. Invisible to machines.