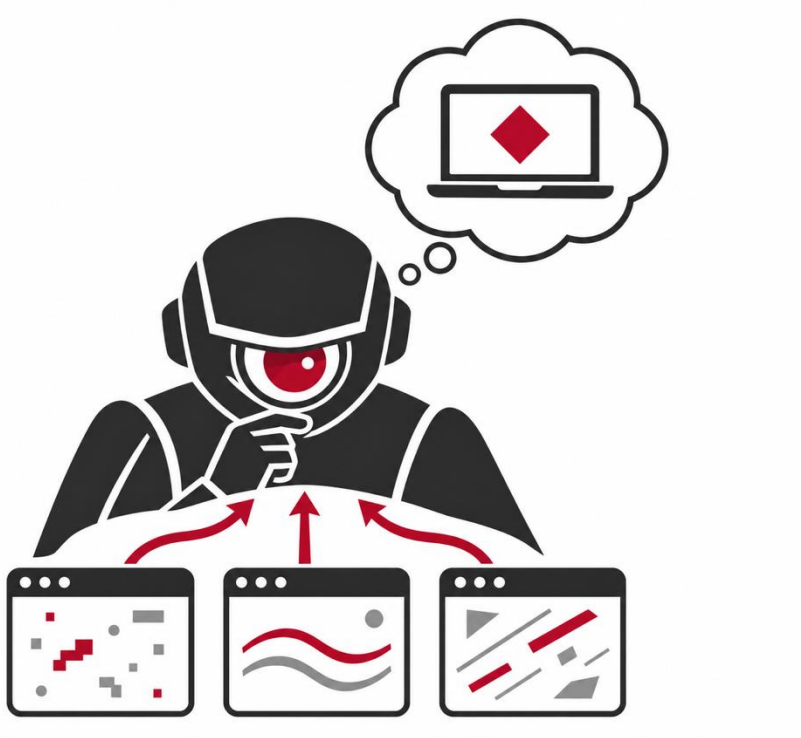
ACM CCS '26 · THE HAGUE

When Tracking Learns

Learning-Based Browser Fingerprinting Under Anti-Fingerprinting Defenses



Xu Lin

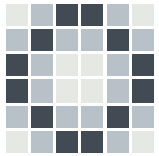
Washington State University

The deterministic story

The tracker



Fixed probes



Same values

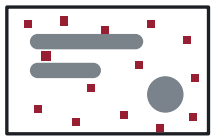
`9f3a2c`

Identifier

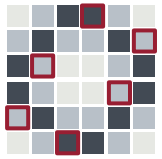
`9f3a2c` = `9f3a2c`

Match

The defense



Add noise



Values change

`e71b04`

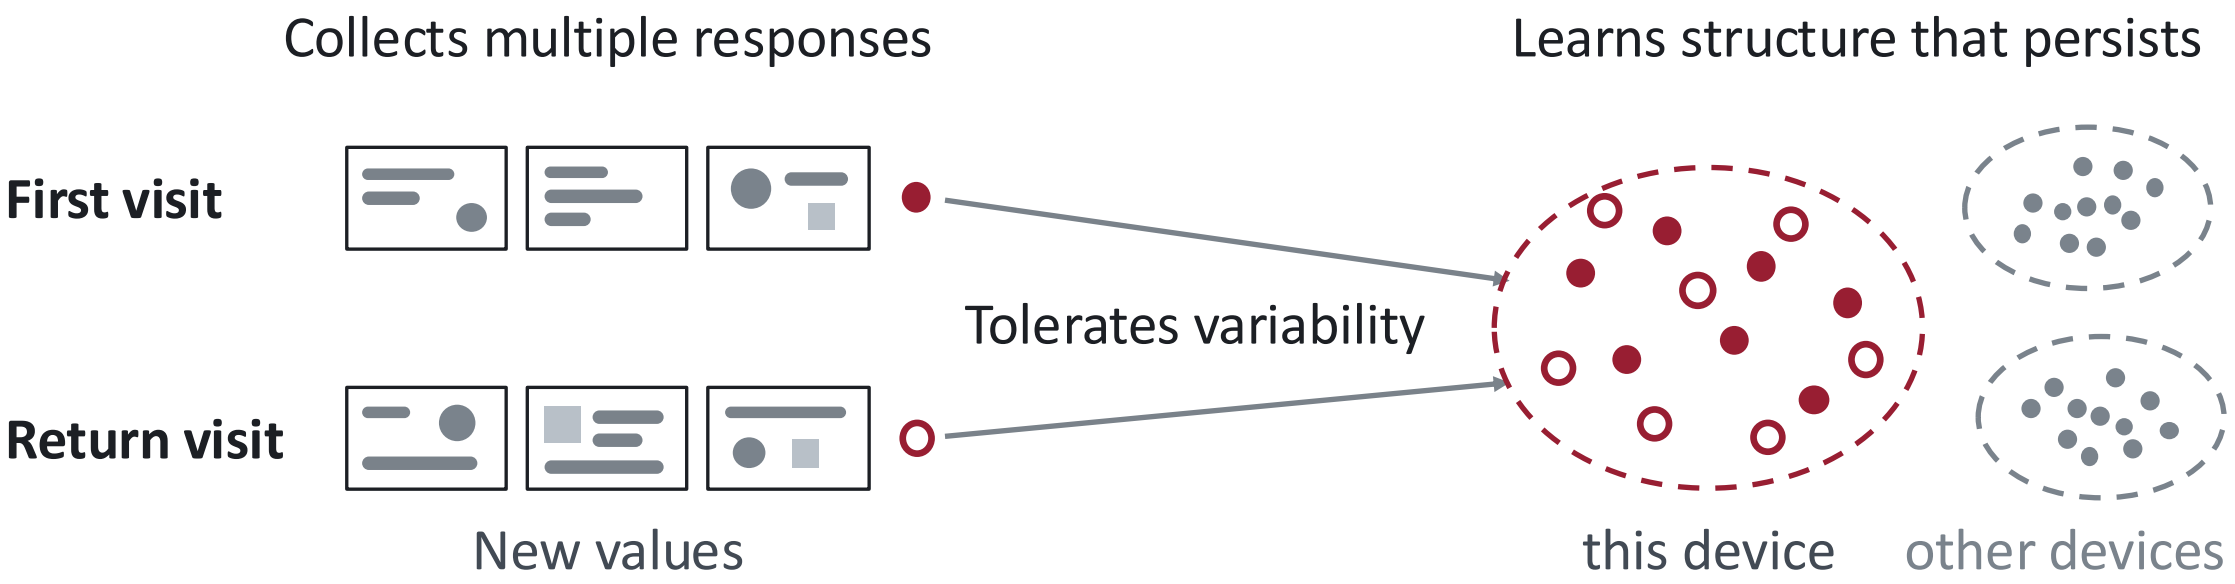
New identifier

`9f3a2c` ≠ `e71b04`

Protected

Defenses are judged by one question: **did the output change?**

What if outputs don’t need to be identical?



New values, same cluster: **the two visits can still be linked.**

A learning-based view

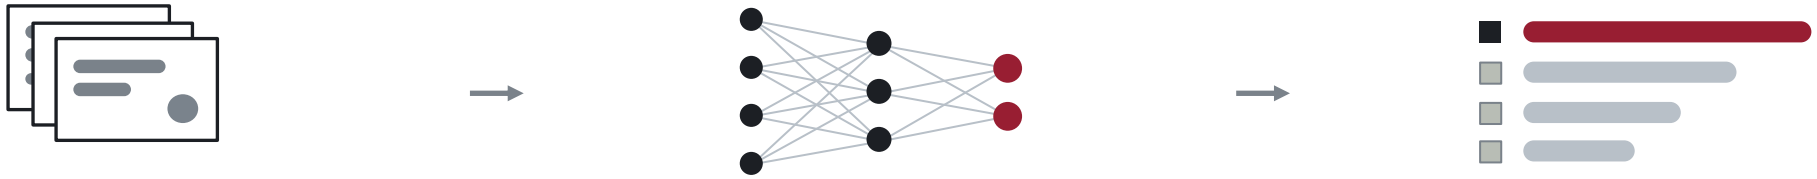
We treat fingerprinting as inference over a device-specific response distribution.

	Exact matching	Learning-based inference
Evidence	one set of values	many responses per device
Linking	the same identifier	similar learned patterns
Variability	can break the match	can carry signal

01 Can we identify devices across visits?

02 Can defenses prevent it?

LearnedFP: a tracker that learns



1 · Collect

Random probes, raw RGBA pixels

2 · Train

One CNN per key, then frozen

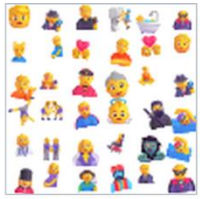
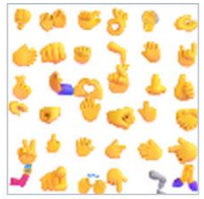
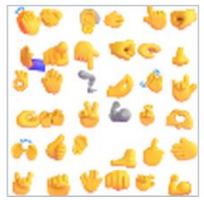
3 · Deploy

Cosine ranking, nine-key fusion

No hashing of browser outputs, and no per-device retraining.

ONE DEVICE, ONE VISIT

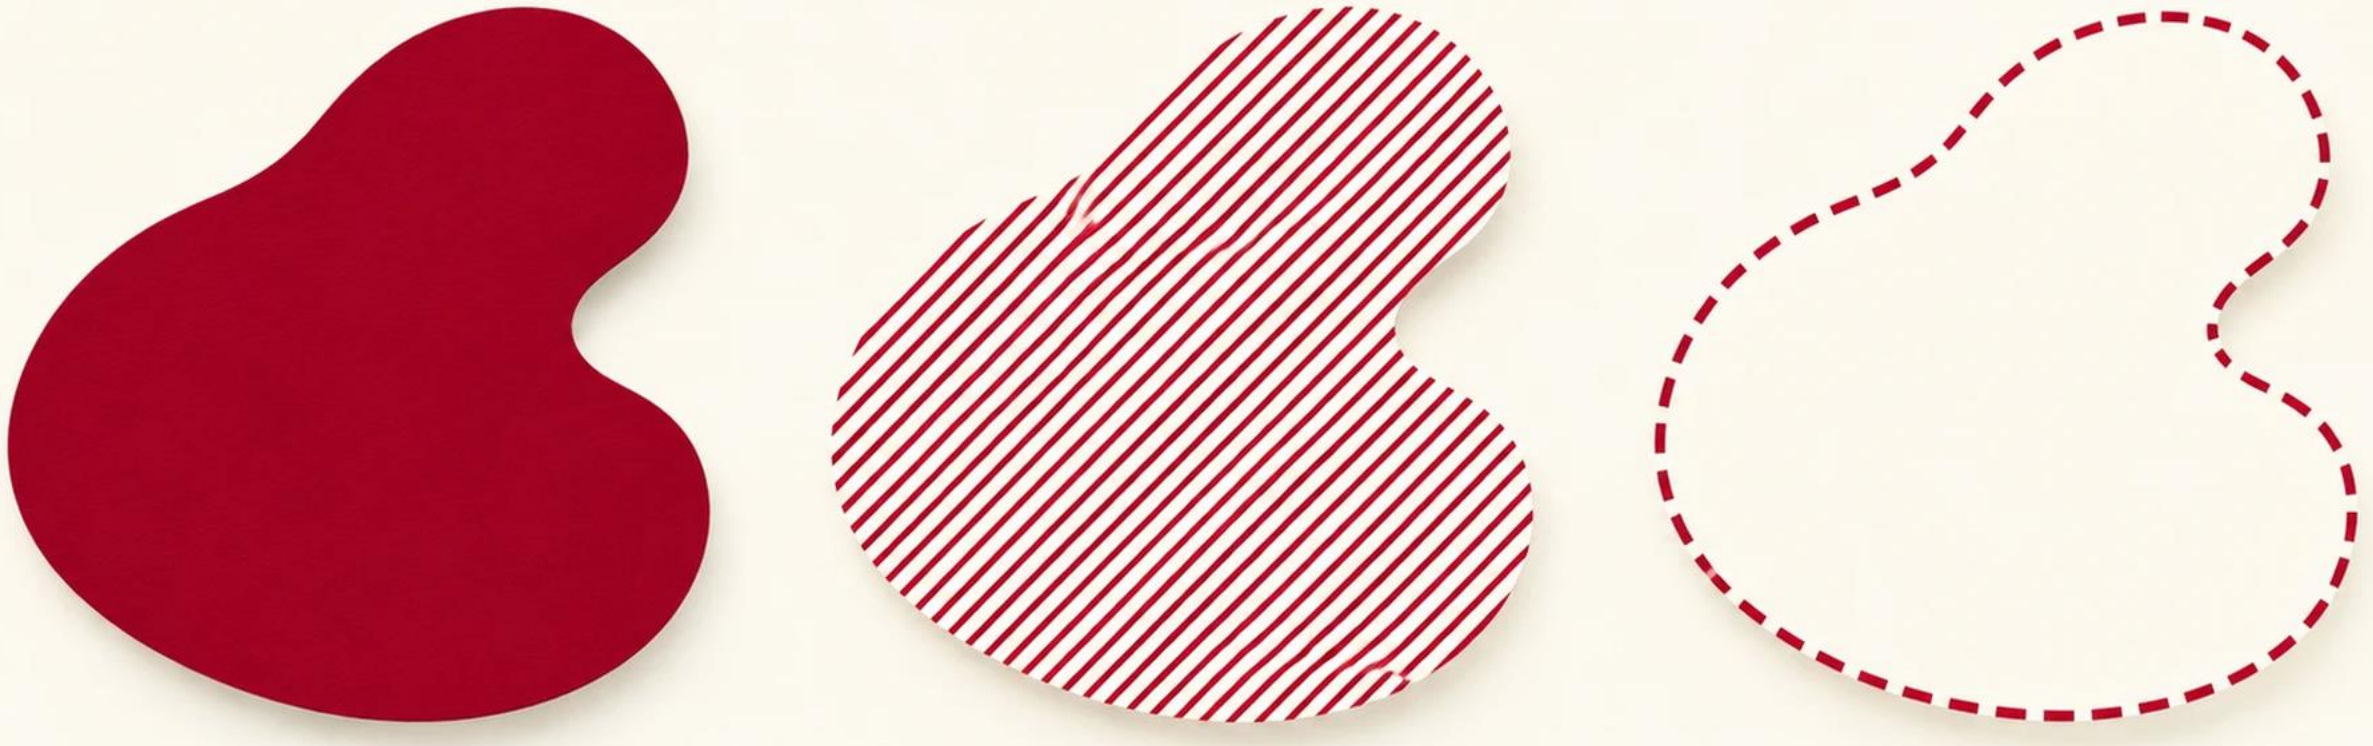
Nine rendering keys

	Emoji				Text		Geometry		
Sample 1									
Sample 2									
	$\approx 10^{530}$	$\approx 10^{544}$	$\approx 10^{549}$	$\approx 10^{510}$	$\approx 10^{43}$	$\approx 10^{272}$	$\approx 10^{32}$	$\approx 10^{26}$	$\approx 10^{53}$

Many variations
within each key

100 samples in ≈10 s
without user interaction

Can we identify devices across visits?



CROWDSOURCED LONGITUDINAL STUDY

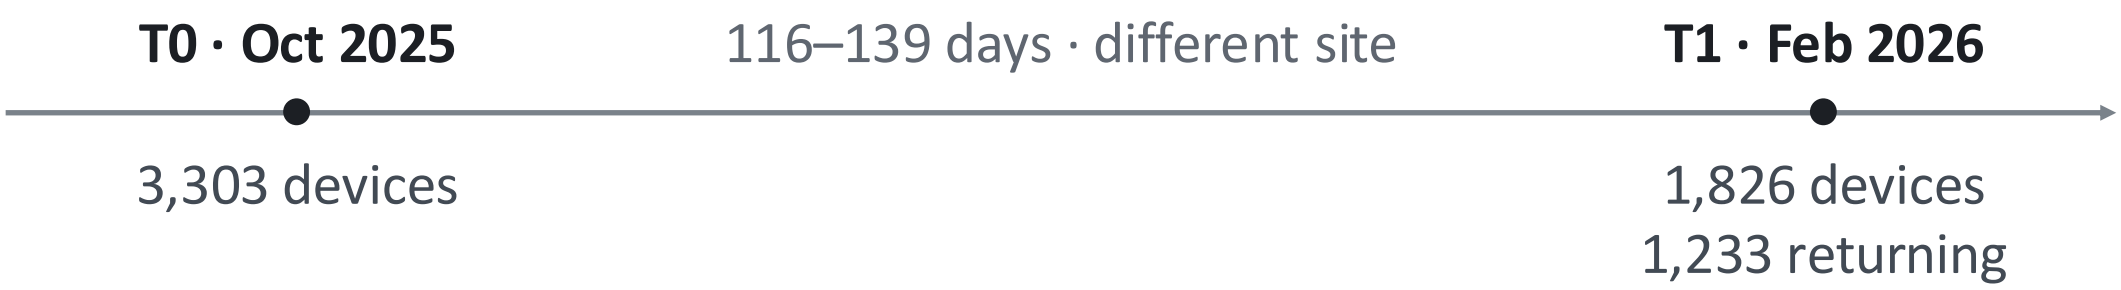
Four months, two sites

3,896

devices from 50 countries

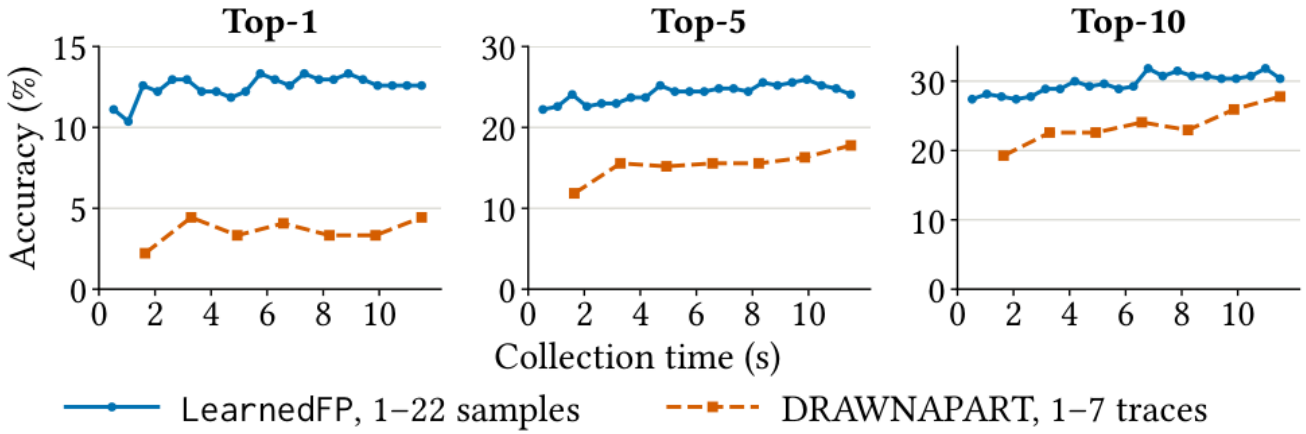
3

systems in the same study:
LearnedFP · DRAWNAPART · FingerprintJS



1,022-DEVICE GALLERY · 270 RETURNING QUERIES

Comparison with the state of the art



0.4 s

first sample: 11.1% Top-1
113× chance

11.5 s

full window: 12.6% vs 4.4% (2.8×)

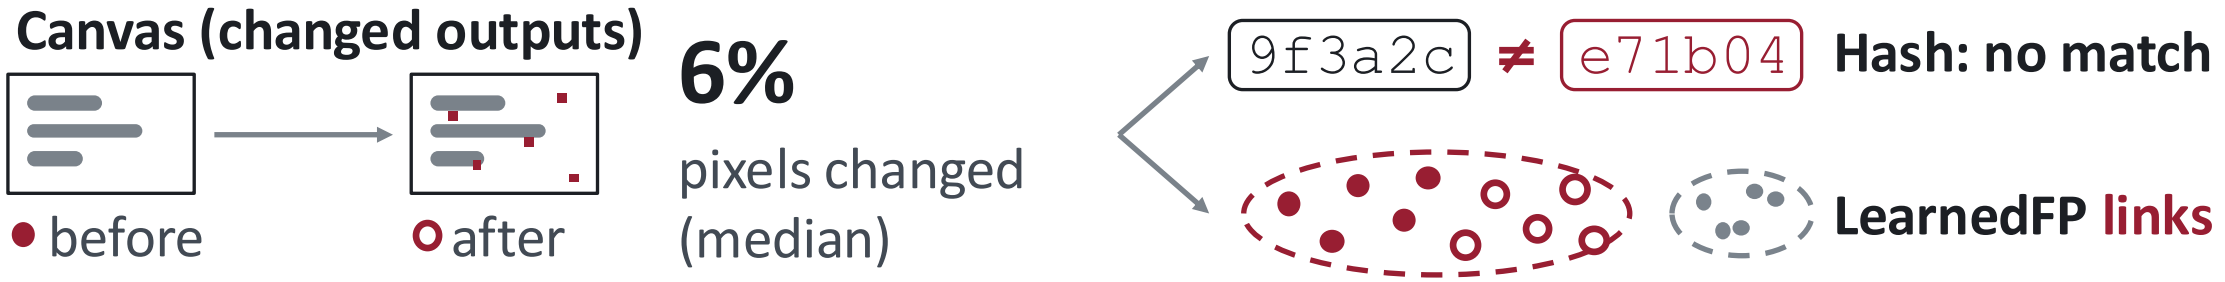
Random baseline: 0.098% Top-1 (1 in 1,022)

Insight 1 Variability is not noise; it is a learnable signal.

FINGERPRINTJS · 3,303 DEVICES · FOUR MONTHS APART

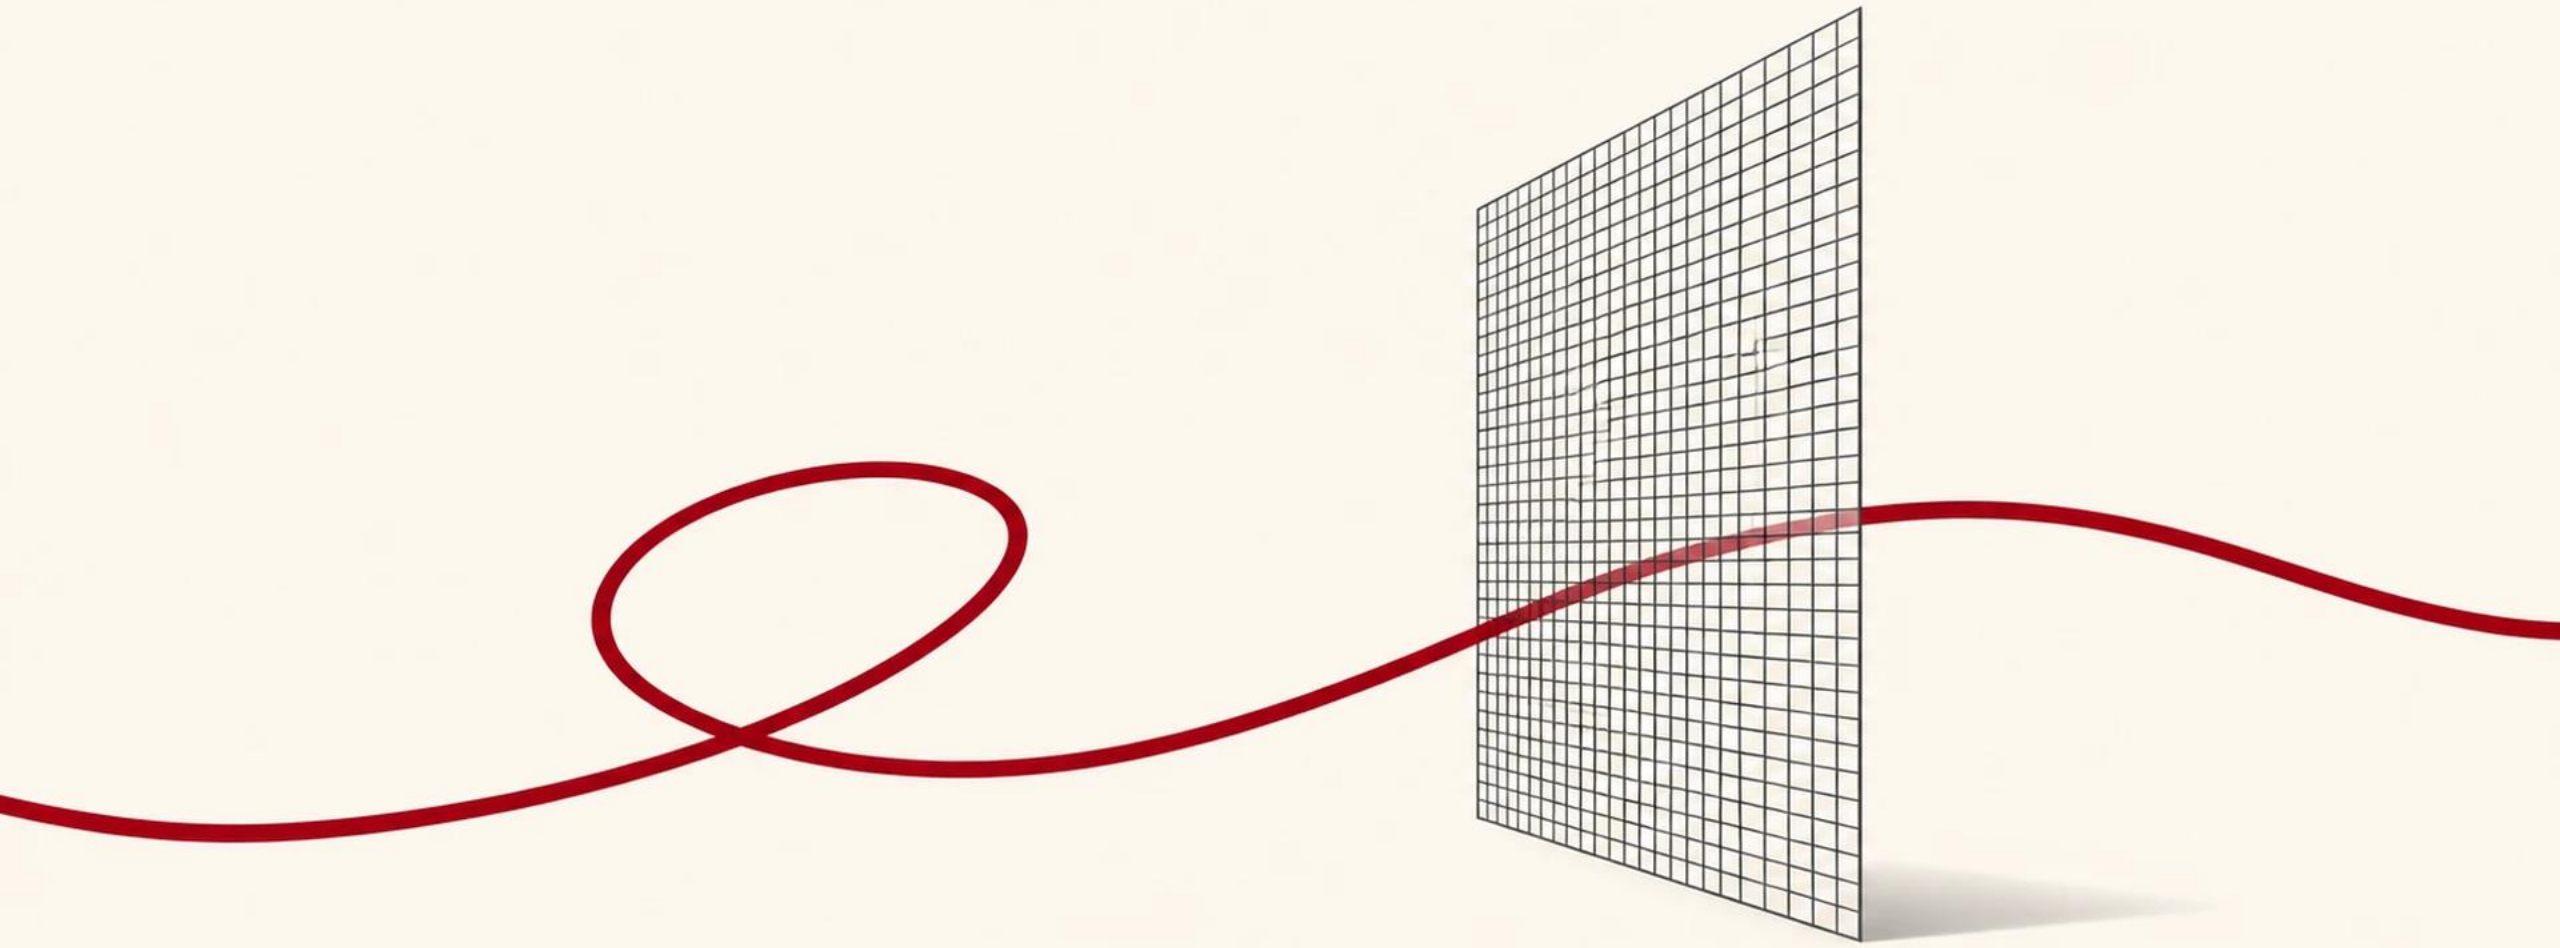
Comparison with exact matching

63% of returning devices have a different identifier



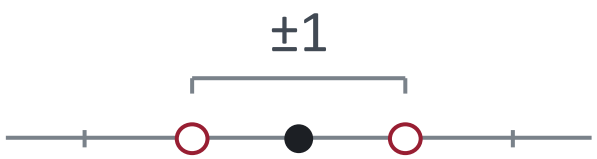
Insight 2 Learning can tolerate small output changes.

Can defenses prevent it?



BRAVE · FIREFOX · SAMSUNG · SAFARI · TOR · 4 EXTENSIONS

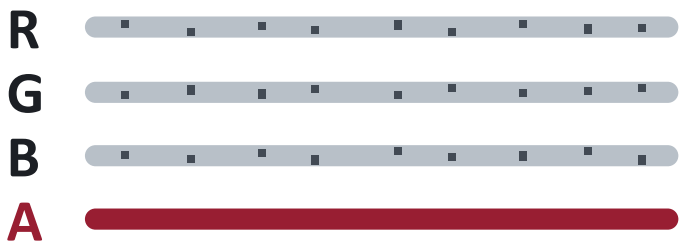
What defenses actually do



Bounded

Brave: Δ ≤ 1 per channel

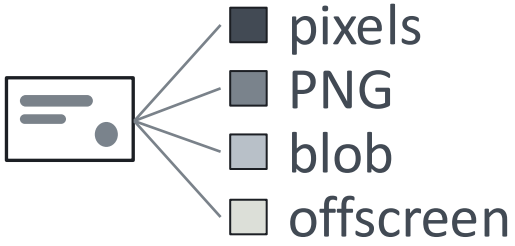
Samsung: ≤ 0.04% of pixels



Partial

Alpha untouched in 5 of 9 defenses

One spoofer misses toBlob()



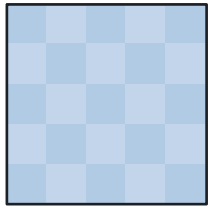
Inconsistent

Extraction paths disagree

Safari: offscreen ≠ direct

LAB SET · 9 DEFENSES · UNDER A SECOND

Every tested defense can be detected



100 × 100 semi-transparent fill → read 1

`rgba(100, 150, 200, 0.5)` → read 2

- **Brave**
  Per-pixel noise, Δ ≤ 2
- **Safari Private**
  Alpha modified
- **CanvasBlocker (kkapsner)**
  Same shift on every read
- **Firefox ETP Strict**
  Noise with Δ > 2
- **Tor**
  Period-8 tile
- **Canvas FP Defender**
  New shift on every read
- **Samsung SAT Strict**
  Only G and B perturbed
- **Canvas Blocker**
  Noise only on a 10×10 grid
- **Fingerprint Spoofer**
  Canvas element swapped

LAB SET · 5 INSTANCES PER CONFIGURATION · 653-ENTRY GALLERY

Defenses become signals



Top-1 · 15 standard vs 45 defended instances

Safari 60% → Private **100%**

Chrome 0% → Canvas Blocker **100%**

Firefox 80% → ETP Strict **60%**

Insight 3 Defenses can become fingerprinting signals.

262 DEVICES · 12 AUDIO GRAPHS

The same recipe on Web Audio



Random audio graph Raw waveform Learned embedding

8.3% Top-1 on return after 5 days (21.7× chance)

Insight 4	The threat is the capability, not the canvas.

Takeaways

	Deterministic view	Once tracking learns
Tracker	matches one stable value	learns from many varied samples
Defense	makes the output change	limits what scripts can extract
Question	Did the output change?	**What can a tracker still learn?**

Thank you!

Questions?

Xu Lin · xu.lin@wsu.edu

Paper



doi.org/10.1145/3830454.3846587

Code

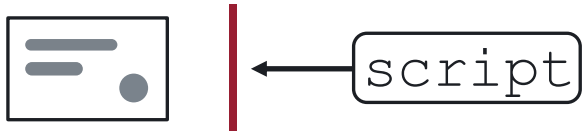


github.com/xlin29/learned-fp

THREE PRINCIPLES · BROWSER EXTENSION

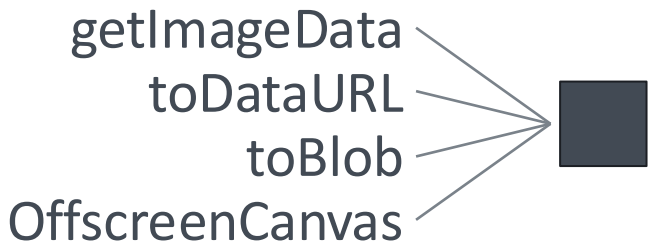
Designing for a learning adversary

1 · Control readback



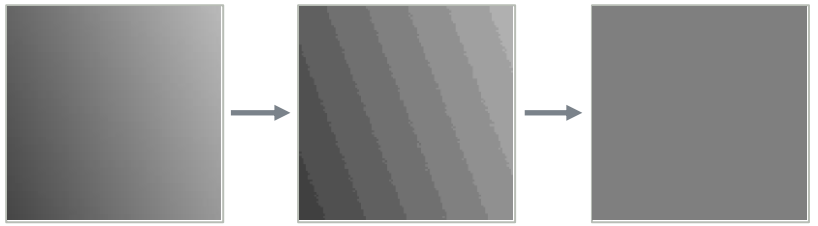
Separate what users see from what scripts read

2 · One snapshot



Same pixels on every path

3 · Budget sampling



Reads 1–3
native

Reads 4–20
4-bit

After 20
constant

Prototype limits: no compatibility study yet · the extension is detectable