


Dictionary Attacks on Biometrics Systems

Nasir Memon

***Tandon School of Engineering
New York University***





NYU CENTER FOR **CYBER SECURITY**

OSIRIS Lab



Computer Security
Club, Hacknights,
Hacker-in-
Residence,
Fellowships

CSAW – Largest In the World



CYBER FELLOWS

- ❖ A unique, affordable online Cybersecurity Master's Degree program designed to address the acute shortage of highly trained, technical professionals in the nation.
- ❖ NYU Tandon is currently offering scholarships of as much as 75% of tuition, bringing the total tuition for the rigorous, highly technical education to approximately \$16,000 for the entire program – the country's most affordable Cybersecurity Master's Degree program by an elite, private university.

Dictionary Attacks on Biometrics Systems

Nasir Memon

New York University




Identity



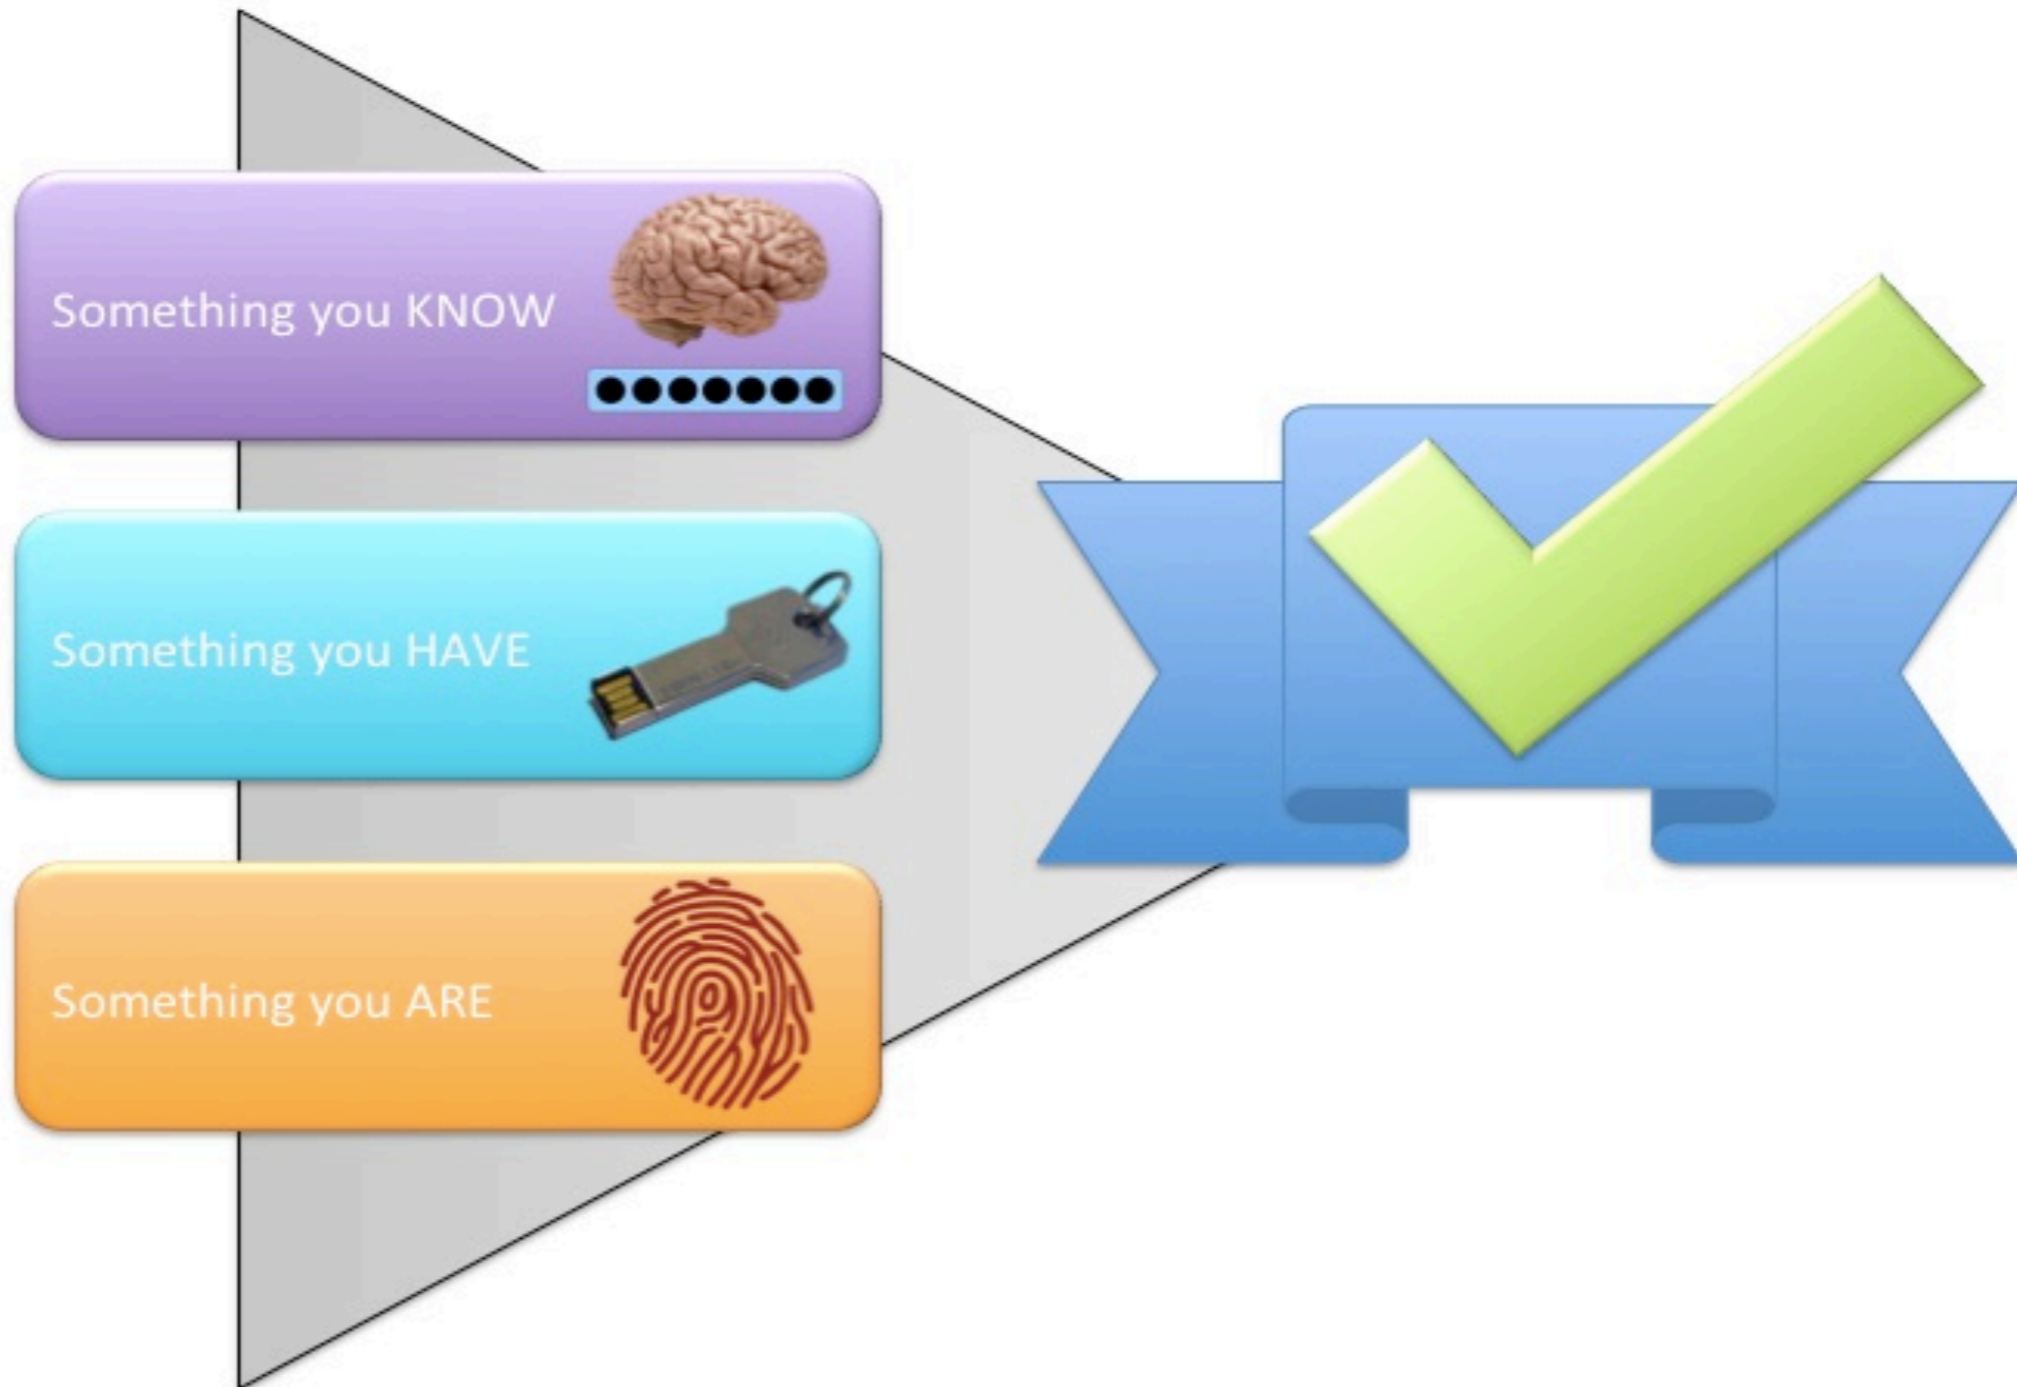
Identity and Authentication

- ❖ What is identity?
 - ❖ A computer's representation of a unique entity (principal).
- ❖ What is authentication?
 - ❖ Binding principal to system's internal representation of identity.
- ❖ Why do we need identity?
 - ❖ Accountability
 - ❖ Access control

Authenticating Computers and Humans





Authentication Approaches




Something-you-know

Usability





Guessing Passwords



SOMETHING YOU ARE - Biometrics



I. Revenues for biometrics industry will more than double between 2007 and 2012 (a). A range of technology modalities will be involved, with the largest share going to conventional and automated and live-scan fingerprinting (b). (courtesy of the International Biometric Group's "Biometrics Market and Industry Report 2007-2012")



Presentation Attacks

- ❖ They exist for many modalities
- ❖ Realistic threat in unsupervised settings
- ❖ Difficult tension between security and convenience
- ❖ Most consumer systems don't have active Presentation Attack Detection systems



Presentation Attacks in Reality

iPhone 5s - Touch ID (Sep 20 2013)

How many days did it take to spoof it ?

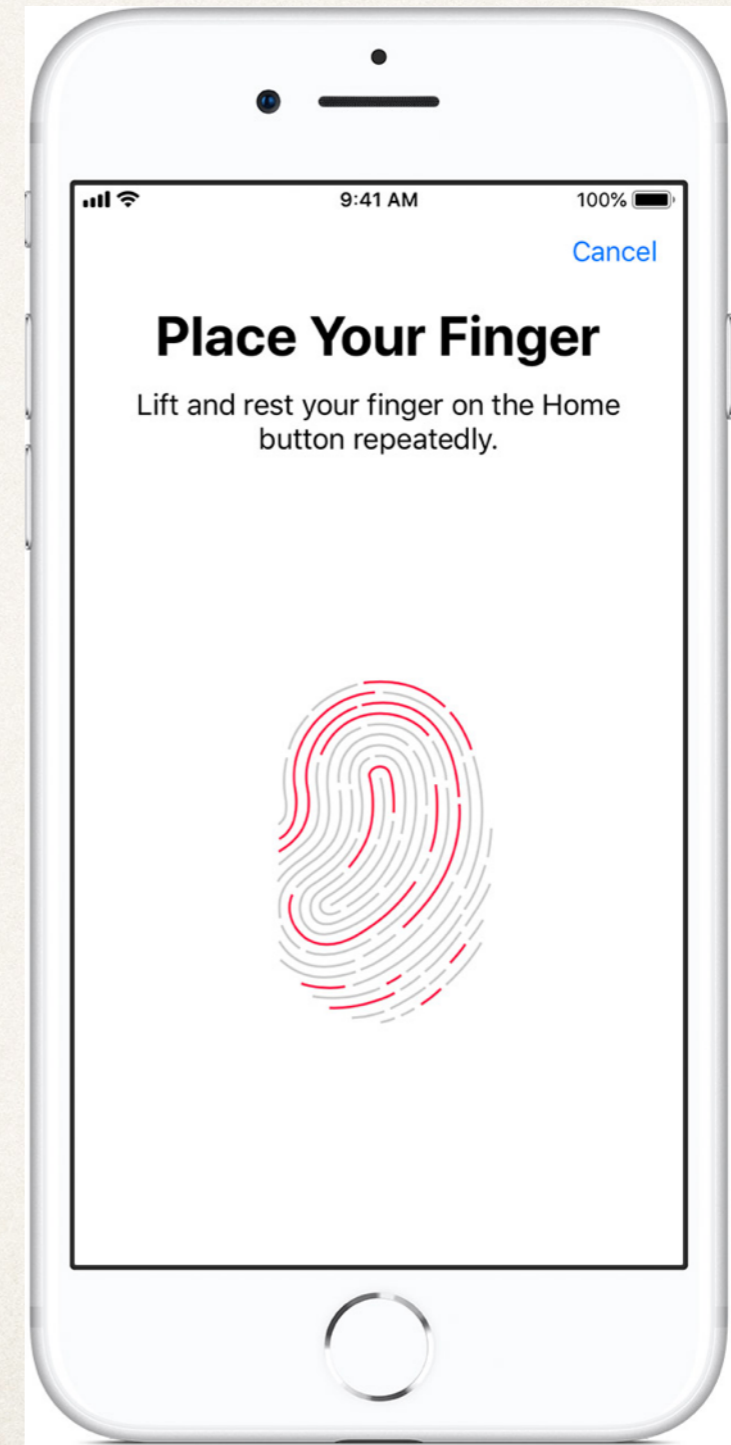
2 days!

- Anti-spoofing techniques getting better
- But so are spoofing techniques
- But in order to spoof, you need victims biometric
- **Does not scale**




Fingerprint Verification on Mobile Devices

- Small sensors
- Enroll multiple fingers



Partial Fingerprint-based Systems

- Capture a limited portion of full finger
- Multiple partial fingerprints are captured
- Which part of which finger is being sensed not known during verification
- Access granted if the sensed partial fingerprint matches any one of the partial fingerprints of any enrolled finger




Biometric Dictionary Attack

- ❖ A type of presentation attack
- ❖ A single biometric that is identified for multiple people
- ❖ It greatly lowers the barrier for an attacker

Rose

Amanda



Mark

Karen

MasterPrints

- ❖ MasterPrints are fingerprints that can **fortuitously** match a large number of other fingerprints [Roy et al. 2017]
- ❖ MasterPrints can be used to launch a **dictionary attack**
- ❖ MasterPrints **maximize the success** of "attacking" a random identity



A. Roy, N. Memon, A. Ross, "MasterPrint: Exploring the Vulnerability of Partial Fingerprint-based Authentication Systems." TIFS 2017

A. Roy, N. Memon, J. Togelius, A. Ross, "Evolutionary Methods for Generating Synthetic MasterPrint Templates: Dictionary Attack in Fingerprint Recognition," ICB 2018

Fingerprint Recognition

- Fingerprints are patterns of epidermal ridges on fingers
- They are highly distinctive (unique) and permanent
- In the last three decades, the focus has been on
 - Cheap and compact sensors
 - Robust and accurate matchers



Fingerprints from the same finger



Fingerprints from two different fingers

Fingerprint Minutiae

- Unordered set of points
- Missing & spurious minutiae, partial fingerprints

Fingerprint

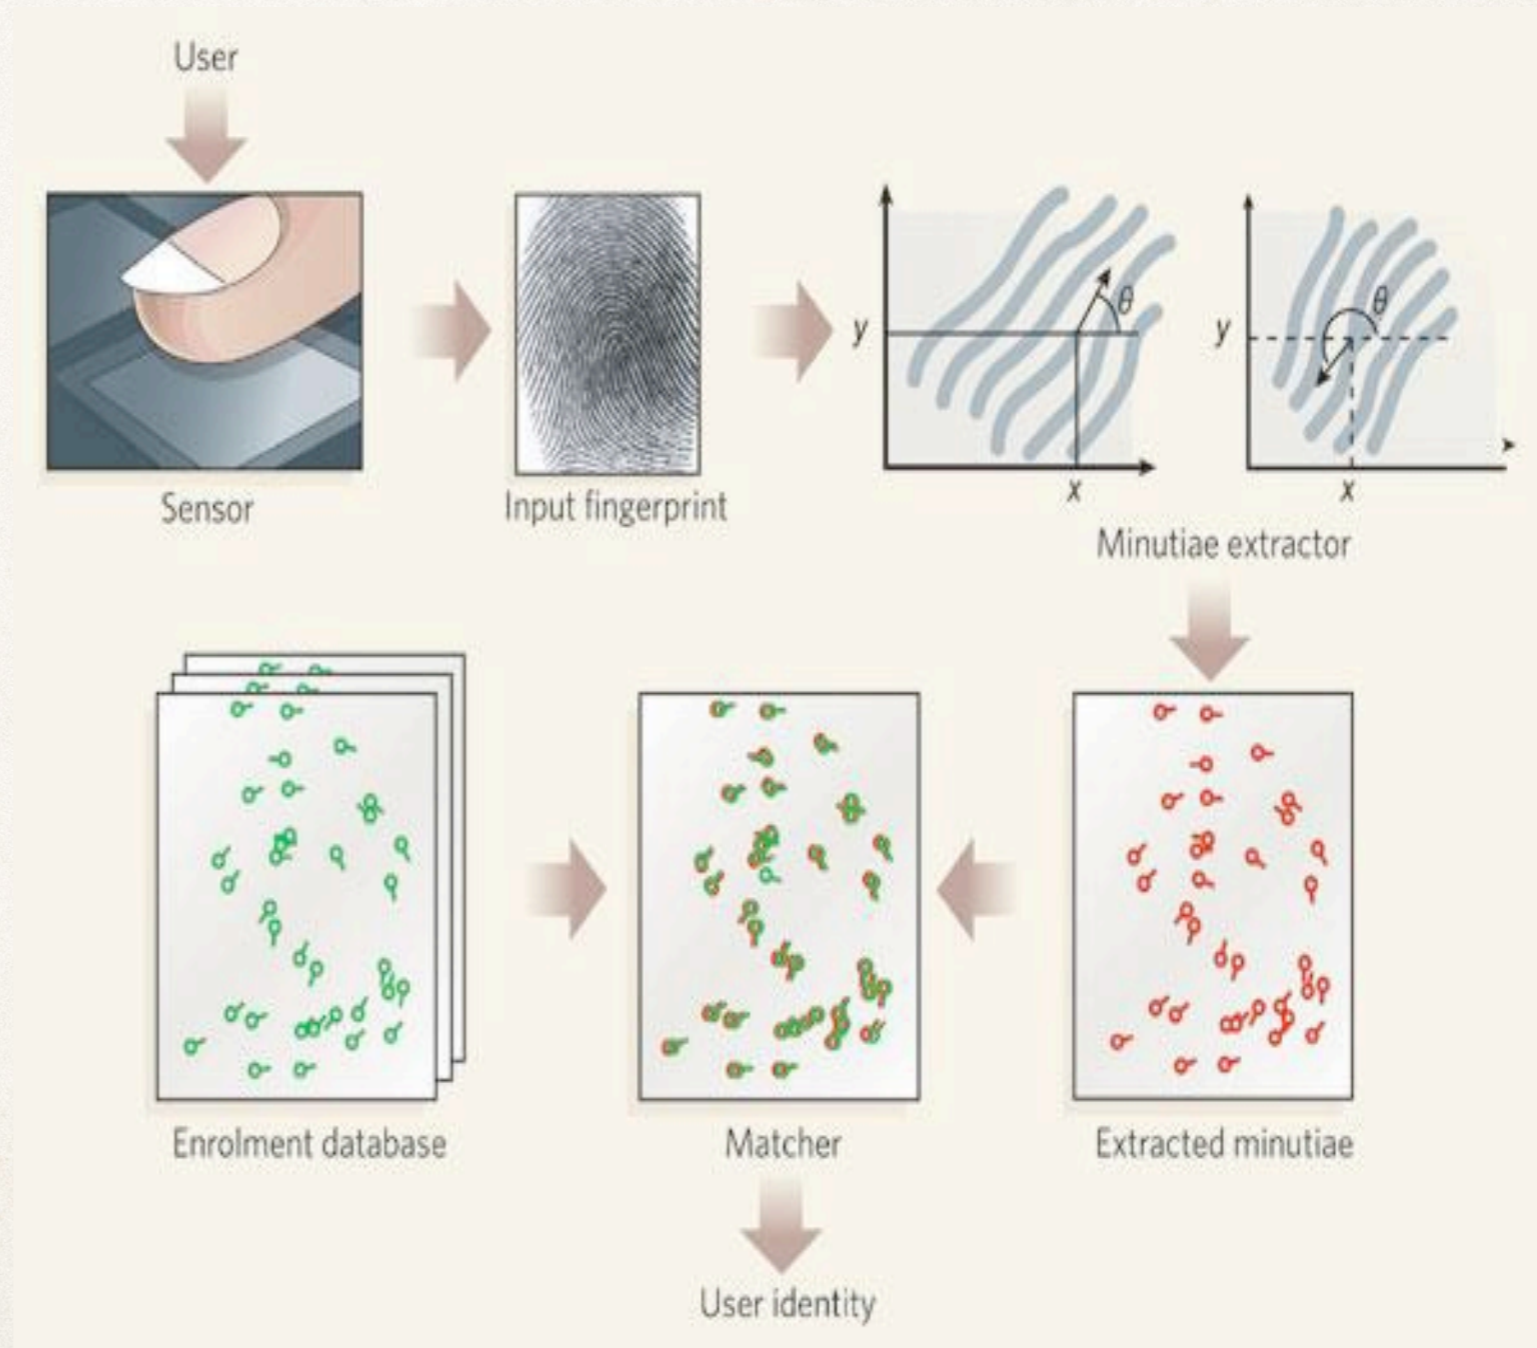


Minutiae (Template)

X	Y	θ
207	138	198
81	144	326
73	158	144
135	203	155
53	205	313

Fingerprint Recognition

Fingerprint Matching:
Find similarity between two fingerprints



Biometric System Errors

- ❖ Important specifications in a biometric system:
 - 1) **FMR**: false match rate (security)
 - 2) **FNMR**: false non-match rate (usability)
 - 3) **FTC**: failure to capture (e.g., a faint fingerprint)
 - 4) **FTE**: failure to enroll

Imposter Match Rate (IMR)

- ❖ **The percentage of false matches when a fingerprint is compared against images of other fingers (impostors)**




MasterPrint Generation

MasterPrints are fingerprints that can **fortuitously** match a large number of other fingerprints [Roy et al. 2017]

- *Sampled MasterPrint (SAMP)* - MasterPrint is sampled from a fixed training dataset
 - *Synthetic MasterPrint (SYMP)* - Generated synthetically
- Both approaches are designed for a minutiae-based fingerprint authentication


Finger Level Comparison with SYMPs

- *Attack against Full Fingerprints*
- SYMPs performed better than corresponding SAMPs
- At 0.1% FMR, IMR increased from 0.4% to 2.8% using 1 impression per finger
- Using 8 impressions per finger, IMR increased from 6.6% to 10.0%




Finger-level Imposter Match Rate using full fingerprint based SAMP and SYMP on the FVC 2002 DB1-A dataset. The IMR of SYMPs was observed to be better than that of SAMPs by ~ 2%.

MasterPrints



Minutiae location of top five partial fingerprints that were selected as MasterPrints from the FingerPass DB7 dataset



Minutiae location of top five partial fingerprints that were selected as MasterPrints from the FVC 2002 DB1-A dataset

Key Observations

- Dictionary attacks are possible using carefully chosen MasterPrints
- With 5 MasterPrints, it was possible to attack 26.46% (FingerPass DB7) and 65.20% users (FVC 2002 DB1-A) at a FMR of 0.1%
- Success varied greatly with FMR and impressions per finger
- Synthetic MasterPrints better than sampled MasterPrints
- High minutiae activity usually occurred in the upper delta point, leading to imposter match with high probability
- Even if a MasterPrint matches a small number of partial fingerprints, the percentage of subjects matched can be quite high
- Risk increases if multiple fingers are enrolled for each subject
- The number as well as the type of partial fingerprints to be stored for each finger should be judiciously chosen to minimize the chance of matching with an arbitrary finger

Questions


- How the distribution of the locations of partial fingerprints affect the attack accuracy?
- Improved synthetic MasterPrint generation technique
- Create synthetic MasterPrint in the “image-level” directly
- How should one select templates with the highest distinctiveness
- Other ways to mitigate the vulnerability associated with the adversarial use of MasterPrints?

DeepMasterPrints

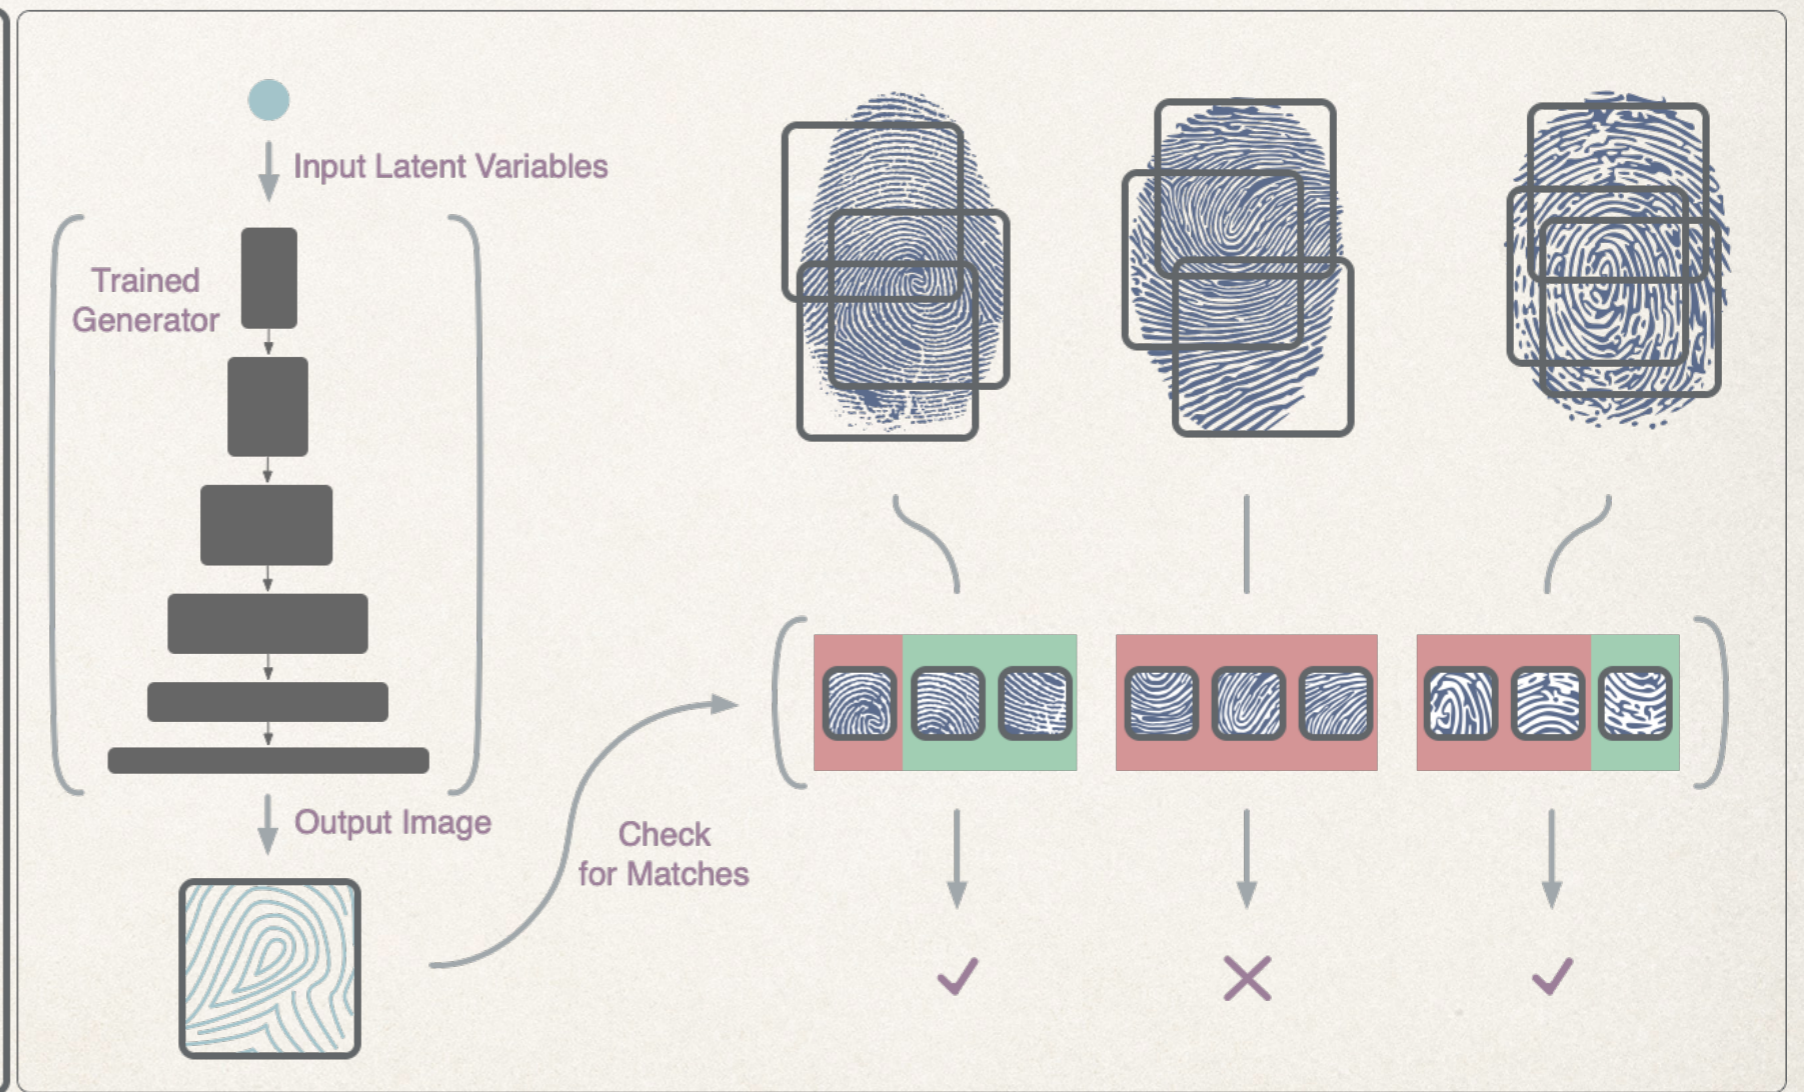
P. Bontrager, A. Roy, J. Togelius, N. Memon, A. Ross, "[DeepMasterPrints: Generating MasterPrints for Dictionary Attacks via Latent Variable Evolution](#)," Proc. of 9th IEEE International Conference on Biometrics: Theory, Applications and Systems (BTAS), (Los Angeles, USA), October 2018. Best Paper Award



Latent Variable Evolution

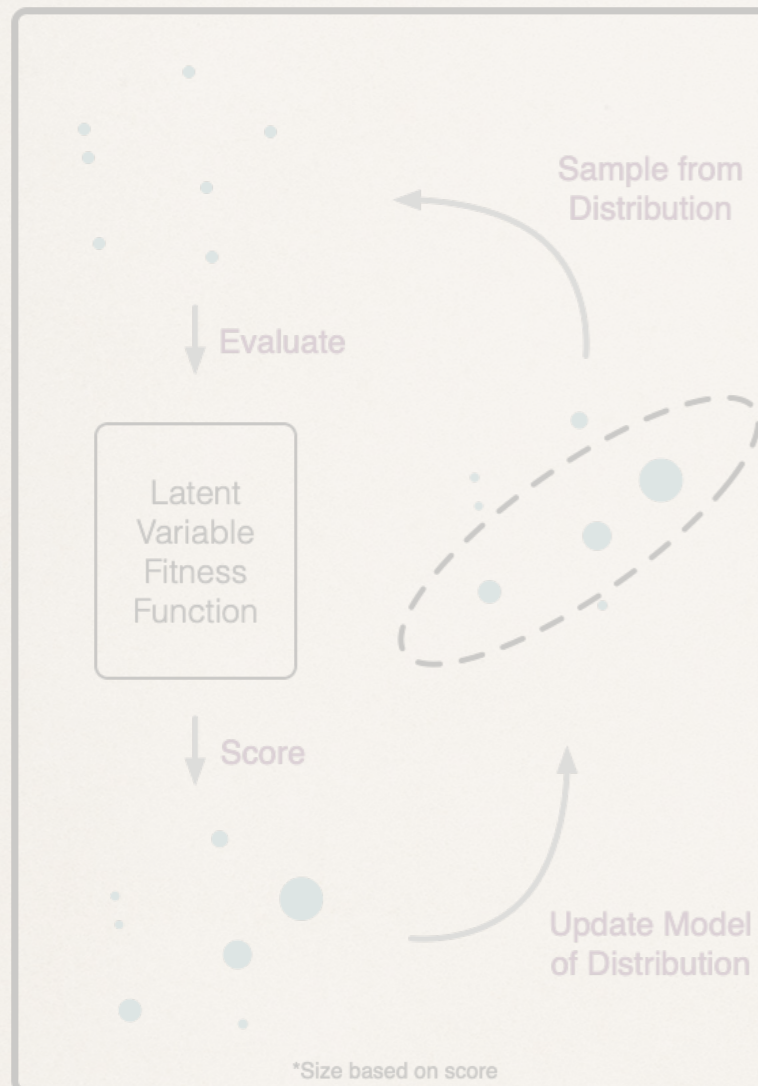


CMA-ES Optimization

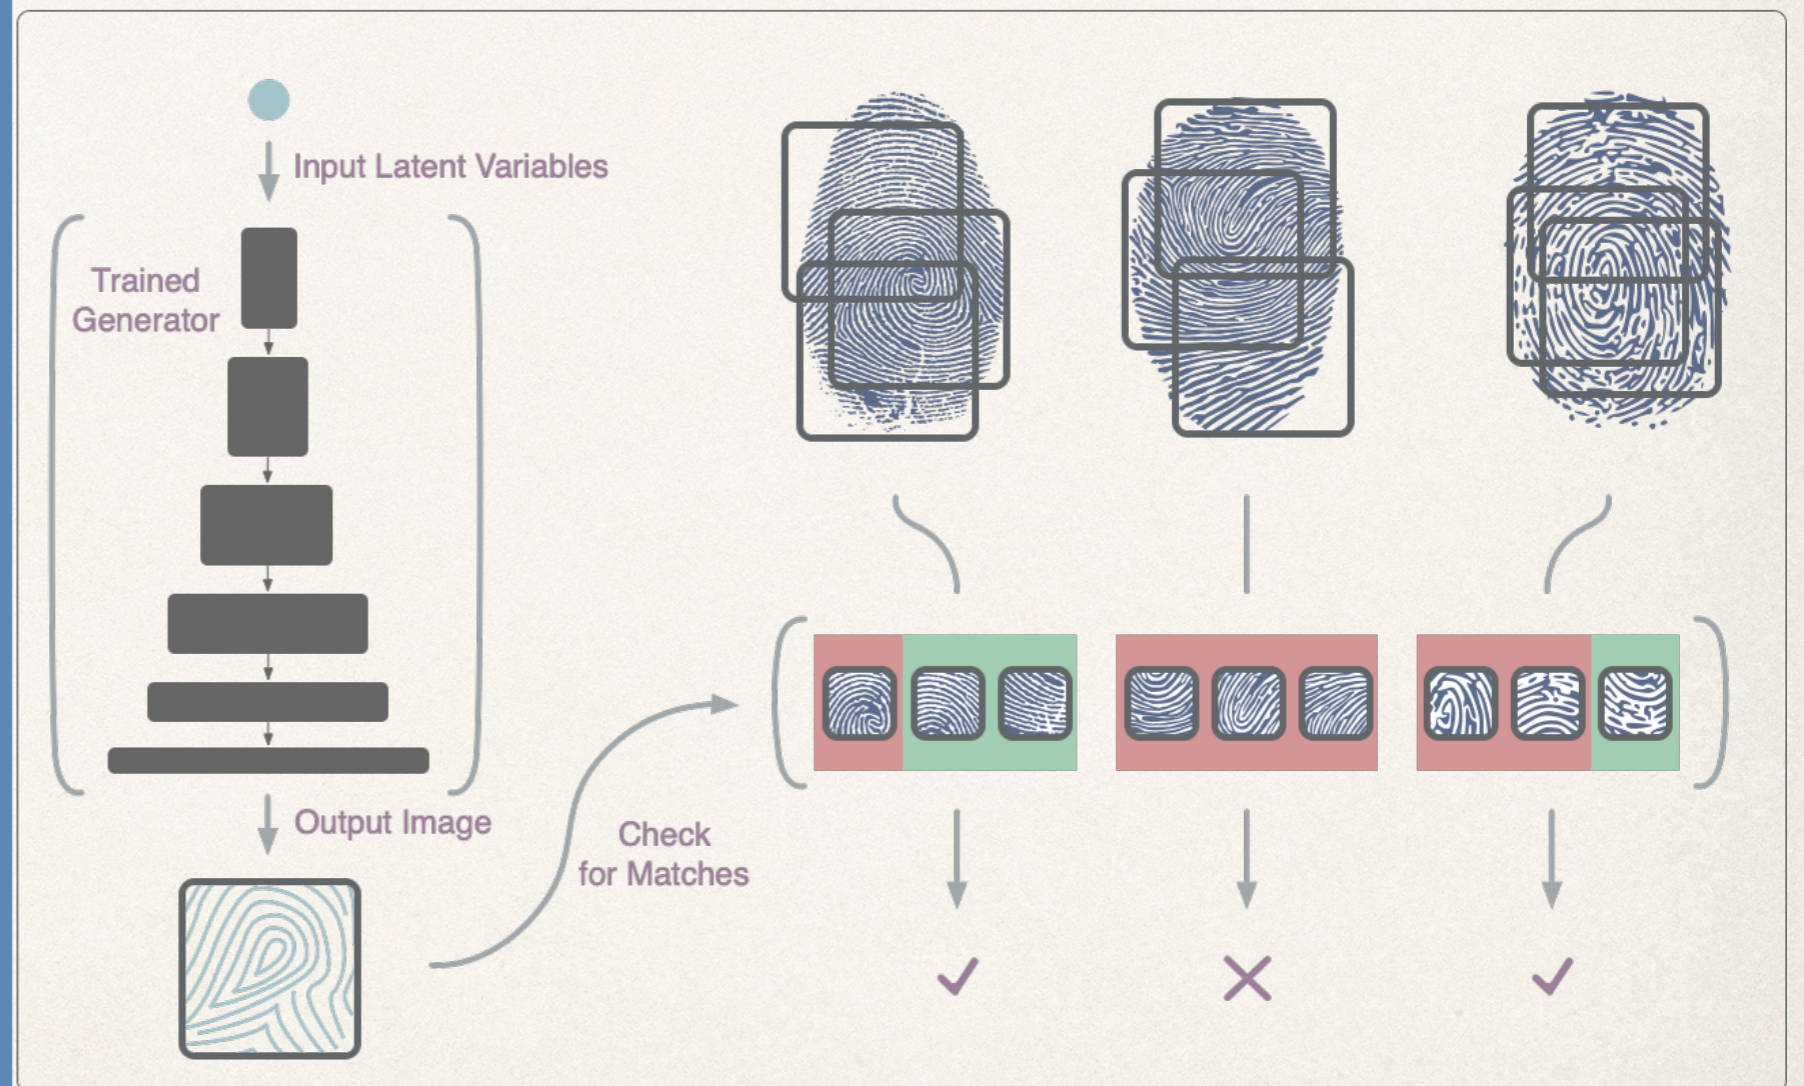


Latent Variable Fitness Function

Latent Variable Evolution

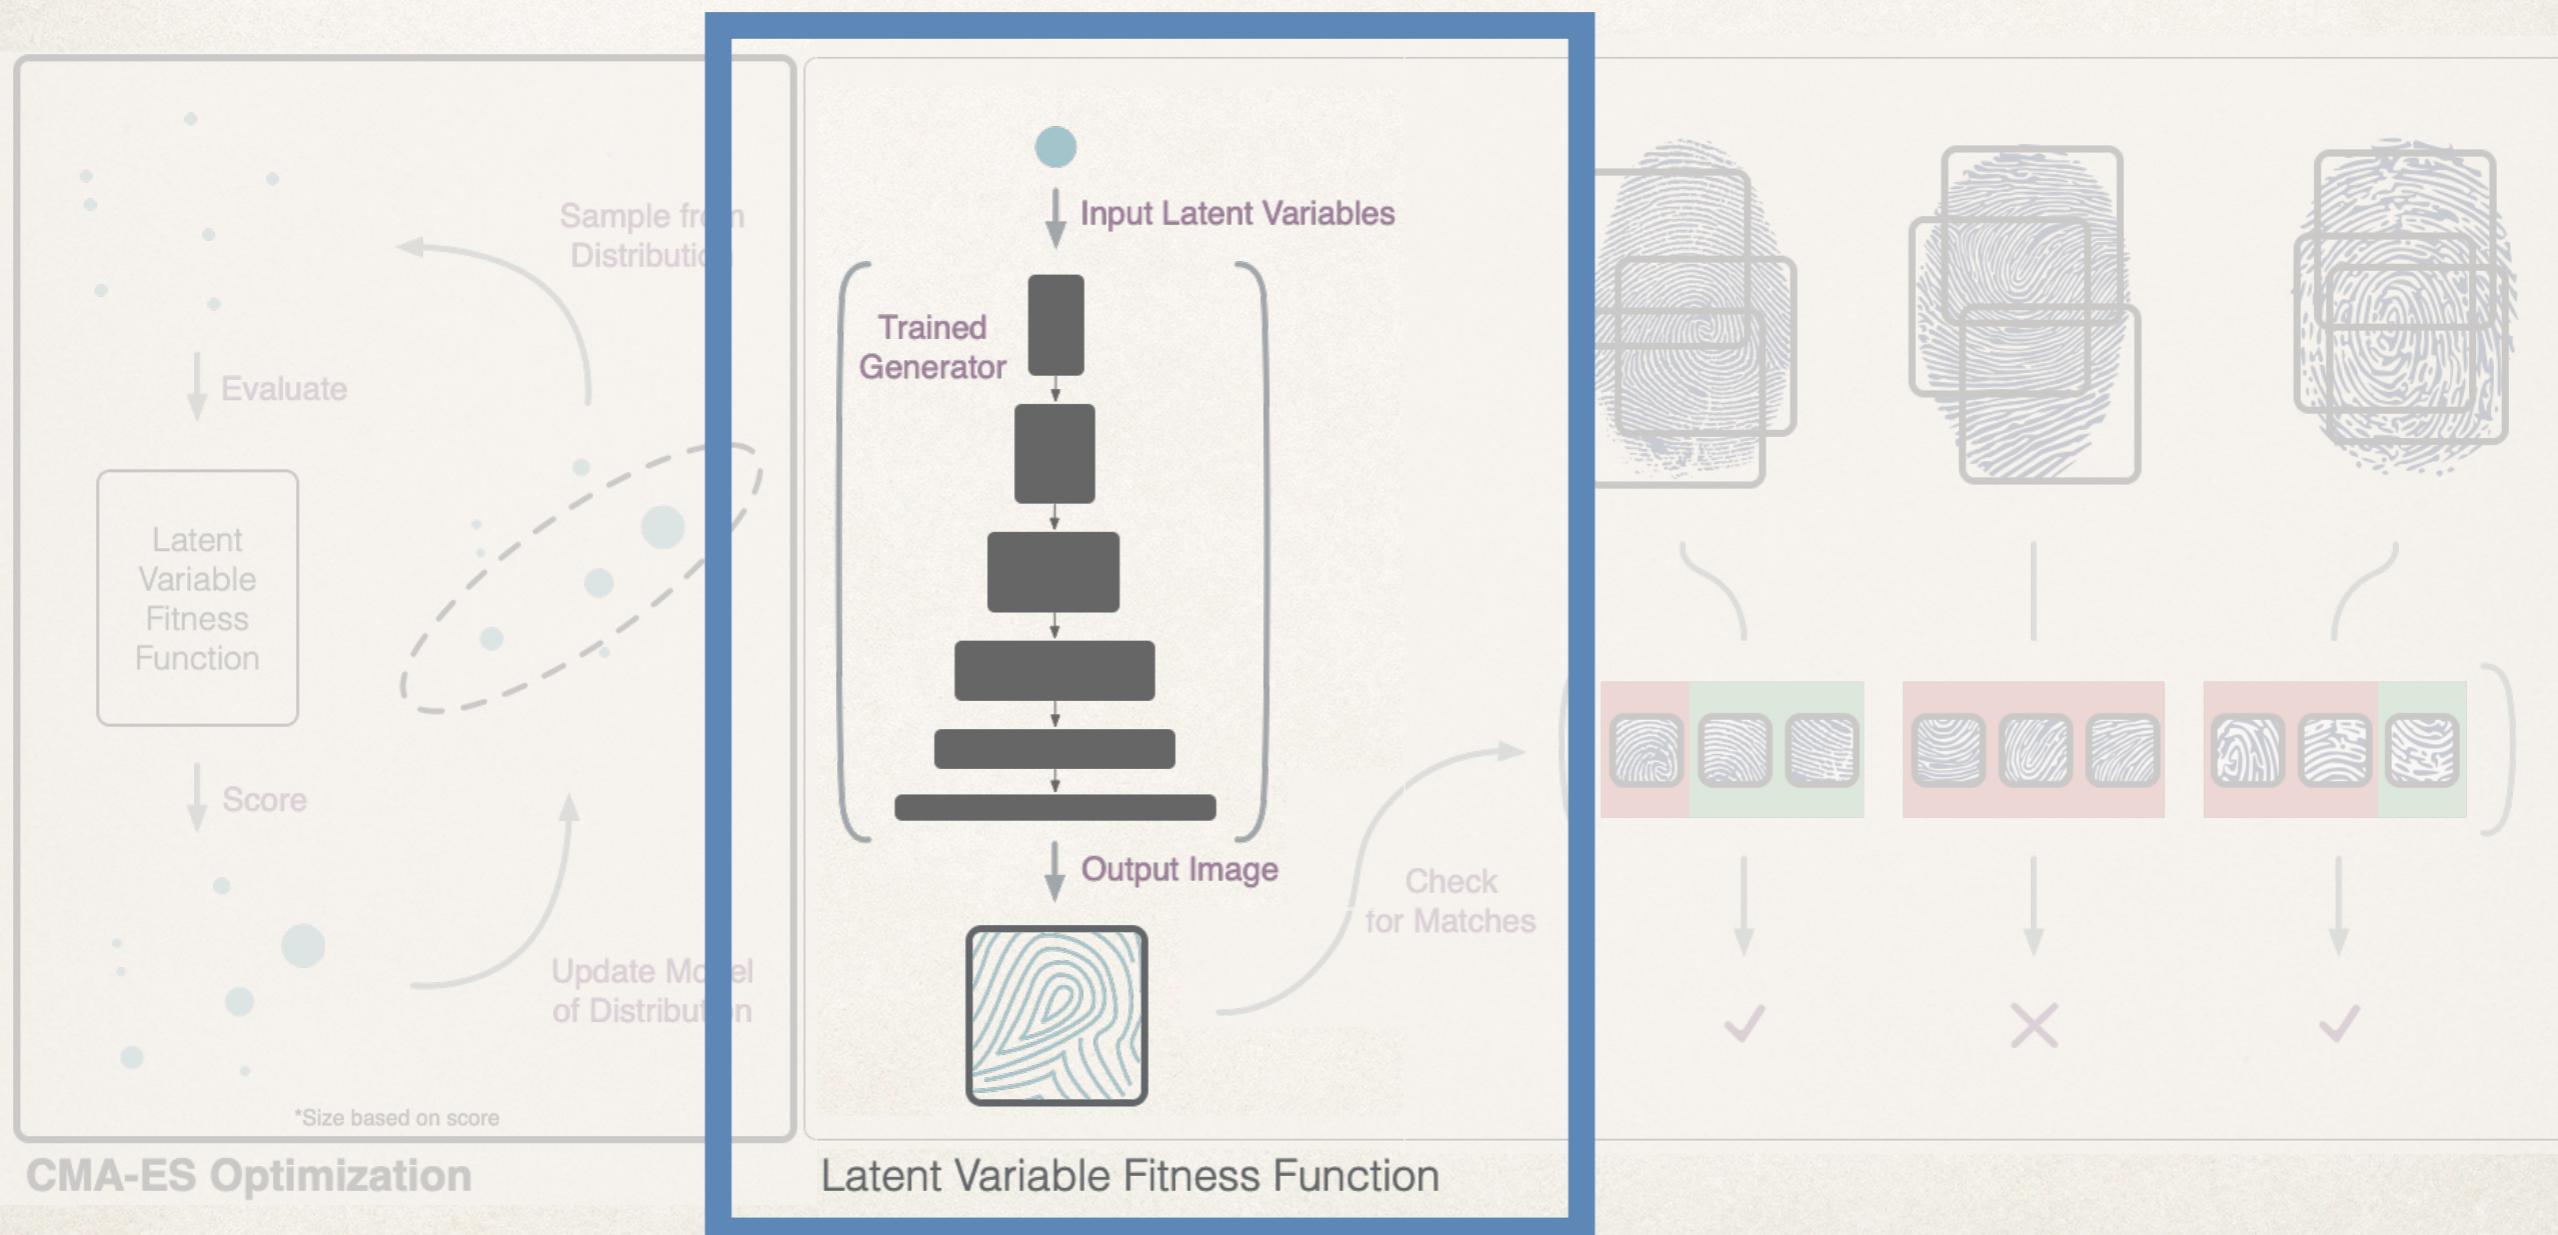


CMA-ES Optimization




Latent Variable Fitness Function

Latent Variable Evolution



Latent Variable to Image

- ❖ Trained Neural Network converts Latent Variables to Images
- ❖ Enforces that all Latent Variables have a resulting fingerprint image
- ❖ Neural Network trained as a Generative Adversarial Network





$$\bullet = [.1, -.4, .4, .3 \dots, .1, 0, -.2, -.3]$$

Latent Variables Sampled from a Normal Distribution

Generative Adversarial Network



Training:




[.1, -.4, ..., -.2, -.3]



Fake




Real




Freeze
Weights

Real




[-.2 .1, ..., .3, .3]

Latent Variable Evolution

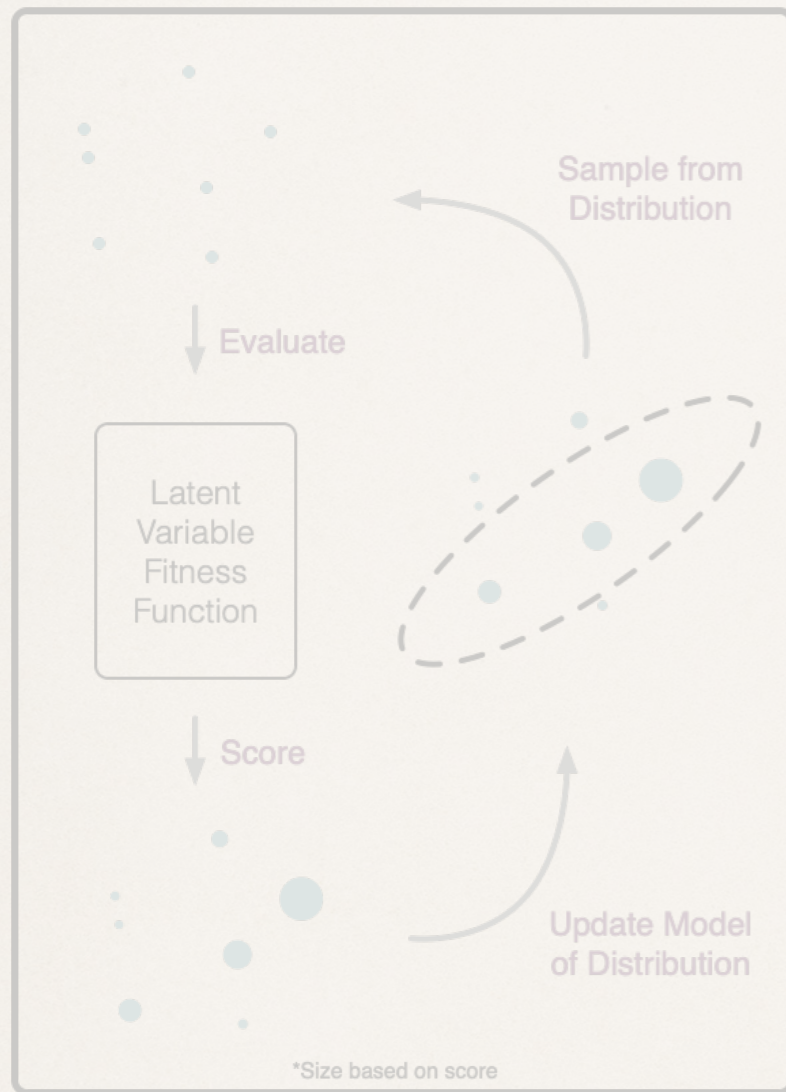


CMA-ES Optimization




Latent Variable Fitness Function

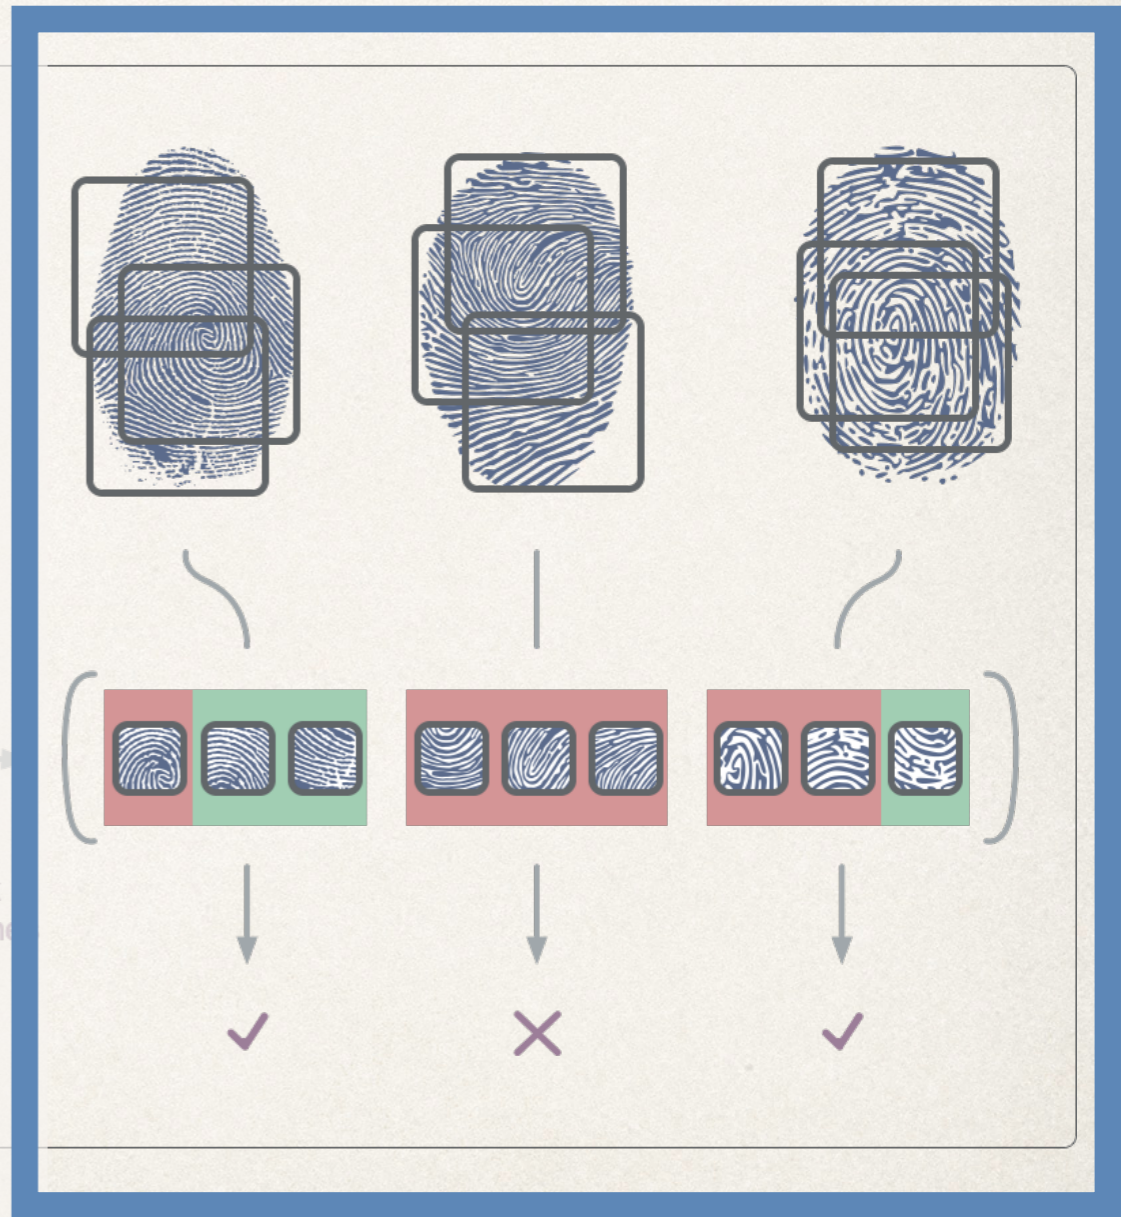
Latent Variable Evolution



CMA-ES Optimization




Latent Variable Fitness Function




Fingerprint Score


- ❖ Test dataset of fingerprints broken apart into partial fingerprints
- ❖ Fingerprint matcher identifies matches at a specific False Match Rate (FMR) threshold
- ❖ A single partial fingerprint match counts as a security breach



Latent Variable Evolution

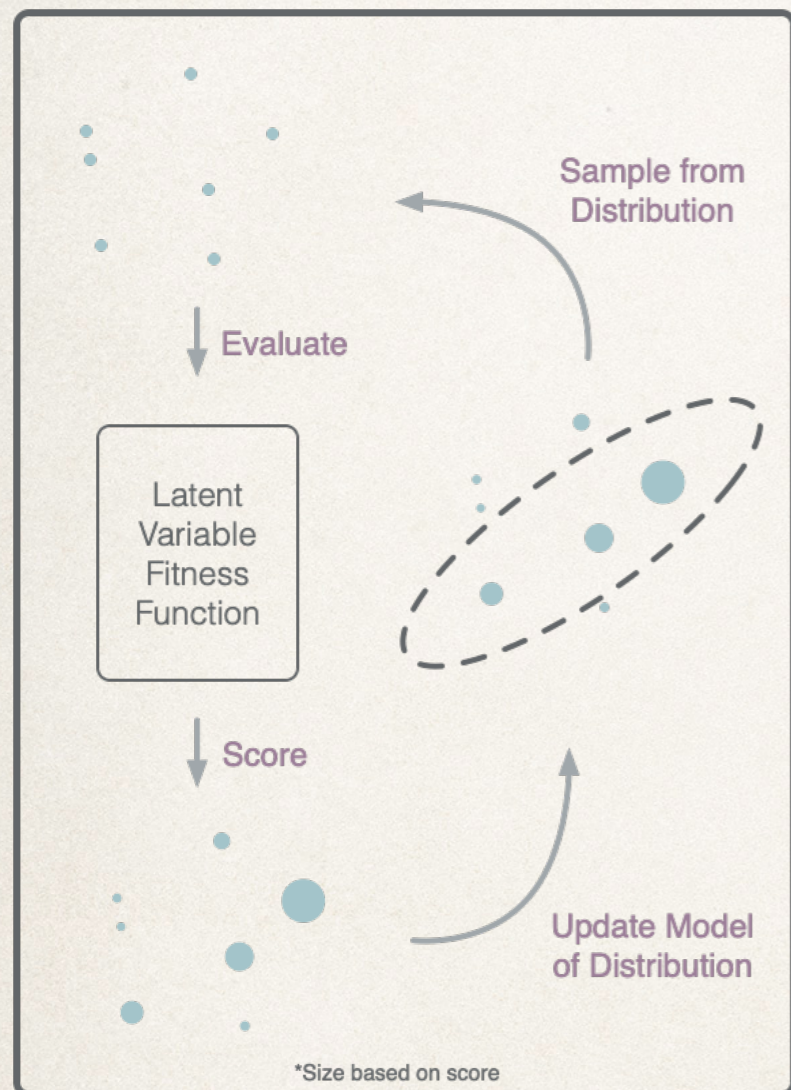


CMA-ES Optimization




Latent Variable Fitness Function

Latent Variable Evolution



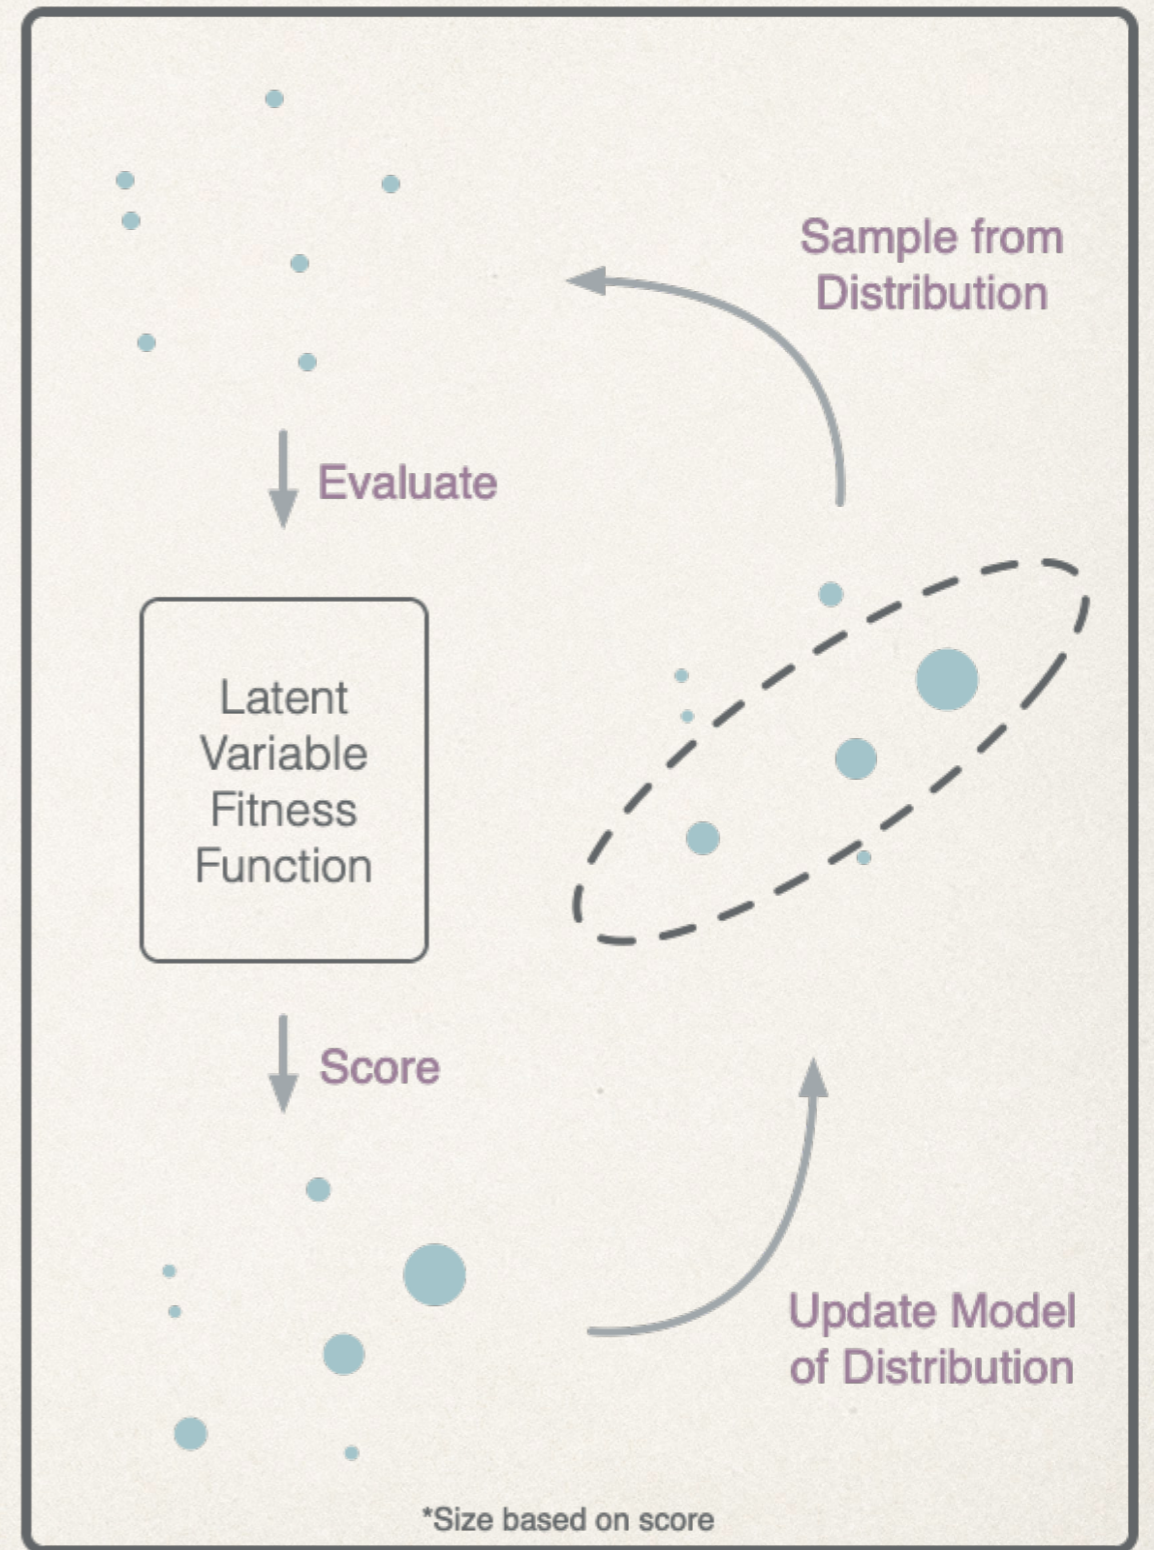
CMA-ES Optimization



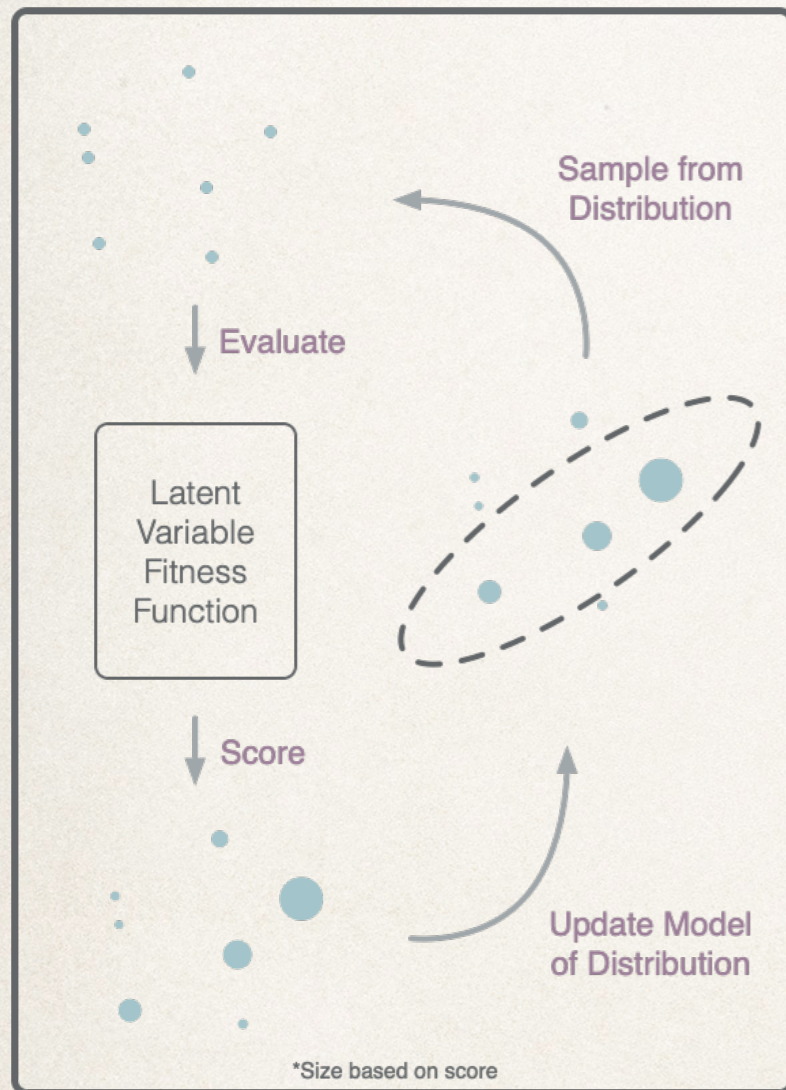
Latent Variable Fitness Function

CMA-ES


- ❖ Sample based evolutionary algorithm
- ❖ DeepMasterPrints represented as latent variables
- ❖ Covariance Matrix of successful fingerprints learned



Latent Variable Evolution



CMA-ES Optimization





Latent Variable Fitness Function

Generated Fingerprints

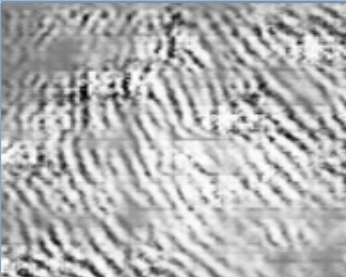
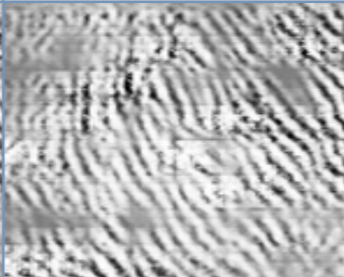
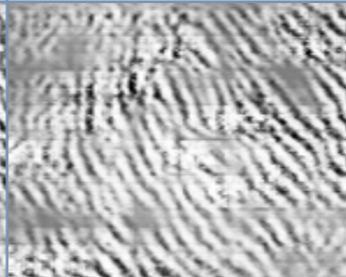

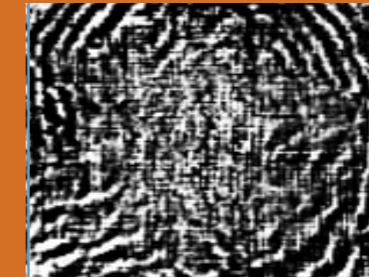

Fingerpass DB7
720 Subjects, Right Thumb



NIST Special Database 9
5400 Subjects, Right Thumb



VeriFinger DeepMasterPrints


	Rolled DeepMasterPrint			Capacitive DeepMasterPrint		
	0.01% FMR	0.1% FMR	1% FMR	0.01% FMR	0.1% FMR	1% FMR
						
Matches	0.3%	8.6%	78.1%	1.1%	22.5%	76.7%

All evolved for the Fingerpass DB7 dataset (50% train / test split)

Comparison to MasterPrints


	0.01% FMR	0.1% FMR	1% FMR
Single MasterPrint	1.9%	6.6%	33.4%
MasterPrint Suite	6.9%	30.7%	77.9%
Single DeepMasterPrint	1.1%	22.5%	76.7%

Master Voices



Mastervoices (MV) - Attacking Speaker Verification

- **Goal:** Attack state-of-the-art speaker verification systems (unconstrained).
- **Output:** Waveform(s) which maximize(s) chance matches with random users.
- **Model:** Pre-trained VGG Vox (contrastive-loss-based VGG).
- **Datasets:** Vox celeb datasets (1 & 2).
- **How:** Use backpropagation to find examples which match many users.



Vox Celeb Datasets

Voice Verification

Training Dataset

VoxCeleb1 - Dev

Users: 1211 (55% M, 45% F)

Utterances: 146,156

Testing Dataset

VoxCeleb1 - Test

Users: 40 (50% M, 50% F)

Utterances: 4,874

Authors' Matlab

VGGVox

7.80 % EER

Our Python VGGVox

8.03 % EER



Master Voice Generation

Exploration/Training Dataset

VoxCeleb2 - Dev

Users: 1000 (50% M, 50% F)

Utterances: 50,000

Testing Dataset


VoxCeleb2 - Test

Users: 1000 (50% M, 50% F)


Utterances: 100,000

* Other 3994 users could be later tested.


Speaker Verification with Vox Datasets & Models




(a) VGGVox Vectors via t-SNE




(b) Female Menagerie




(c) Male Menagerie




(d) Per-Gender Utterance FAR




(e) Utterance Rank per Female



(f) Utterance Rank per Male

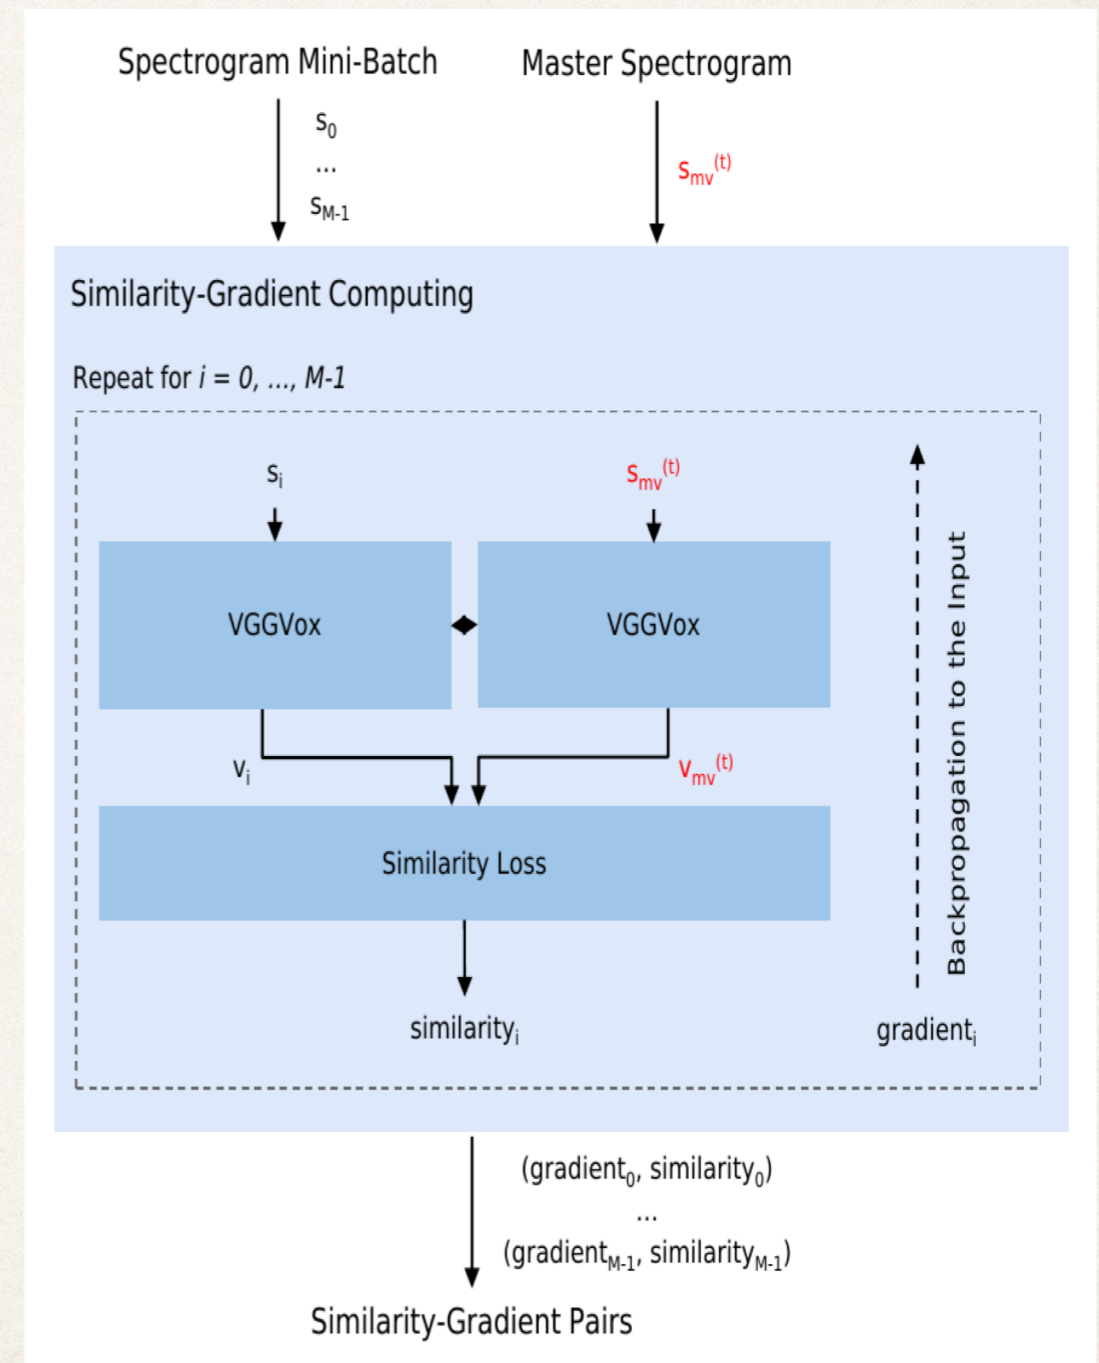



(g) Impersonation Rate@Any10

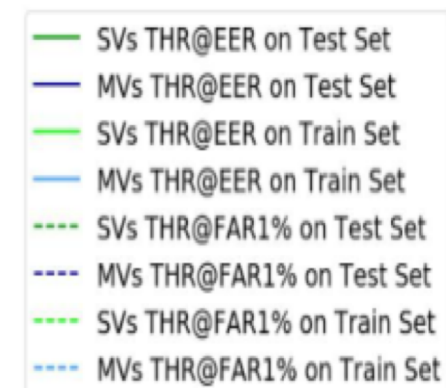




(h) ImpersonationRate@Avg10

Attack Protocol

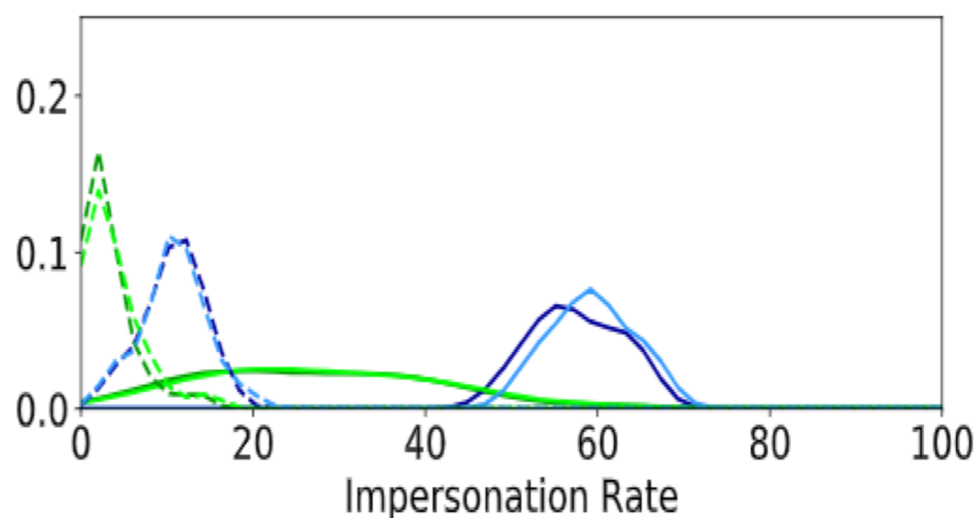



Inherent & Optimized Impersonation Rates



(a) *F2F Impersonation Rate Distribution @ Any10*

(b) *M2M Impersonation Rate Distribution @ Any10*



(c) *F2F Impersonation Rate Distribution @ Avg10*

(d) *M2M Impersonation Rate Distribution @ Avg10*

Conclusions

- Human speech seems susceptible to dictionary attacks.
- SOTA speaker verification models are **not suitable for security applications**:
 - Low TPR with reasonable FPR constraints,
 - High FPR with reasonable TPR levels,
 - Huge differences between male and female speakers (data or model problem?).
- Adversarial attacks allow to quickly improve false matching rates for arbitrary voice samples.
- MV **transfer well to a different population.**
- MV **are robust to spectrogram computation & inversion.**
- Matching strategies can have huge security implications.

Ongoing Work

- Better selection of templates for partial fingerprint scenario
- Real playback attacks.
- Attacking deployed speaker verification in voice assistants.
- Other modalities
- Narrower targeting