Archives available at journals.mriindia.com

International Journal of Advanced Electrical and Electronics Engineering

ISSN: 2278-8948

Volume 13 Issue 01, 2024

An IoT Face-Recognition Door-Lock System

[1]Anusha S, [2]Babitha Rajilin R.C, [3]Jeba Esther Babitha S, [4]Vinushiya N, [5]A. Annie Steffy Beula

[1,2,3,4,5]Department of Electrical and Electronics Engineering, Stella Mary's College of Engineering, Aruthenganvilai, Kanyakumari District, Tamil Nadu 629202, India

Peer Review Information

Submission: 11 Feb 2024
Revision: 07 March 2024
Acceptance: 05 April 2024

Keywords

Face Recognition, Door Lock, Home Automation, YOLOv4, Object Detection, Arduino, False Accept Rate, Presentation Attack Detection, Biometric Access Control, Internet of Things.
Note on Sources for This Revision

Abstract

This paper reports the design of a face-recognition door-lock system for home automation, in which a webcam at the entrance feeds a computer running a YOLOv4 detector, a set of enrolled face images is held as a database, and an Arduino Uno drives the relay that releases the lock when a presented face matches the database. The paper formalises the detection, matching and decision relations the system rests on, and then examines what a lock of this kind has to be measured by. Three findings follow, and each is quantitative. First, YOLOv4 is an object detector: it answers whether a face is present and where, and does not answer whose face it is. Recognition additionally requires alignment, an embedding and a matching step against the enrolled set, and the algorithm performing that step is not named in the project record. Second, a door lock is a biometric access-control system and its security is the false accept rate at a stated threshold, not an accuracy percentage. Because the lock performs open-set identification against every enrolled person rather than verification against one, the system false accept rate grows with the number enrolled: a per-comparison rate of 0.1 per cent gives a system rate of about 1 per cent with ten residents enrolled and about 5 per cent with fifty. Converted into operational terms, at ten impostor presentations per day a system rate of 1 per cent admits a stranger about once every ten days. Third, a webcam-based system without liveness detection is defeated by a printed photograph, and no anti-spoofing stage appears in the described design. The paper also notes that the recognition decision is taken on the computer and conveyed to the Arduino as a serial command, so the serial link is the trust boundary, and that the behaviour on computer failure is unspecified. A four-trial measurement plan is given. This paper was prepared from the project abstract alone, so status entries are marked for confirmation rather than asserted.

Nomenclature

Symbol	Meaning
e, e_i	Embedding of the presented face and of enrolled identity i
s	Similarity between two embeddings
τ	Matching threshold at which access is granted
N	Number of enrolled identities
FAR, FRR	False accept rate and false reject rate at a stated threshold

FAR_{sys}	System false accept rate over all enrolled identities
EER	Equal error rate, where the two errors are equal
r	Impostor presentations per day
T_{int}	Mean interval between false accepts (days)
T_{e2e}	Time from face presented to lock released (s)
t_{cap}, t_{det}	Frame capture and detection time (s)
t_{emb}, t_{match}	Embedding and matching time (s)
t_{act}	Actuation time of the relay and lock (s)
AP	Average precision of a detector

Introduction

Home automation has made connected devices inexpensive, and access control is among the first functions householders want automated. A lock that opens to a recognised face is attractive because it removes the two failure modes of a key and of a password: a key can be lost or copied, and a password can be forgotten or shared. A face is carried by its owner and cannot be mislaid [1], [2].

It is also, however, a credential with properties that a key does not have. It cannot be changed after a compromise; it is presented in public and can be photographed; and a system that reads it makes a probabilistic decision rather than a deterministic one. A face lock therefore has to be characterised in a way a mechanical lock does not, and that characterisation is the subject of this paper.

The system described here places a webcam at the entrance, runs a YOLOv4 detector on a computer to locate faces in the video, holds a set of enrolled face images as a database, and uses an Arduino Uno to drive the lock relay when a presented face matches. This paper documents the design and sets out the figures such a system must be measured by before it could be relied on.

1. Problem Statement

Reported face-recognition lock prototypes are numerous and nearly all report the same evidence: the door opened for an enrolled person and did not open for someone else. That is a functional demonstration, not a security assessment. A biometric access-control system is characterised by its error rates at a stated operating threshold, by its behaviour as the enrolled population grows, and by its resistance to a deliberately presented artefact. These are the quantities that determine whether a lock is safe to fit, and they are rarely reported. This paper makes them explicit and computes what they imply.

2. Objectives

The specific objectives of the work are as follows:

- To build a door-lock system that detects faces at the entrance, matches them against an enrolled database, and releases the lock only on a match.
- To separate detection from recognition, and to identify precisely which stage the named algorithm performs.
- To state the error-rate relations by which a biometric lock is judged, and to show how the system error rate depends on the number of people enrolled.
- To convert those rates into operational terms a householder can act on.
- To identify the presentation-attack and trust-boundary exposures created by this architecture.
- To specify the trials that would produce the figures the design currently lacks.

3. Scope and Delimitation

The study covers the system design, the detection and matching relations, and the specification of the trials a lock claim requires. It reports no accuracy, false accept rate, false reject rate, latency or spoof-resistance figure, because this paper was prepared from the project abstract and no such data were available to it. It makes no claim about what the campaign recorded. It also makes no claim of suitability for fitting to an occupied dwelling, for the reasons given in Section VI.

4. Contributions

1. A clear separation of detection from recognition, identifying which stage YOLOv4 performs and which stages the record leaves unspecified.
2. A statement of the error-rate relations by which a biometric lock is judged, in place of the single accuracy percentage such reports usually offer.
3. A computed demonstration that the system false accept rate grows with the number of enrolled identities, because the lock performs open-set identification rather than verification.
4. A conversion of that rate into an operational interval, giving the expected time between admissions of a stranger.

5. Identification of the presentation-attack exposure and of the serial link as the architecture's trust boundary.
6. A four-trial measurement plan producing the figures a lock claim requires.

The remainder of the paper is organised as follows. Section II reviews related work and fixes the published detector figures. Section III develops the relations. Section IV describes the system. Section V examines what the lock claim requires. Section VI treats the security and safety questions, Section VII the limitations. Section VIII gives the measurement framework, and Sections IX and X the conclusion and the future scope.

Related Work and Reference Figures

Automated face recognition for door access has an established literature, and the systems reported range from classical subspace methods to modern deep embeddings [1], [2]. The methodological point common to the serious work in the field is that a recognition system is reported at an operating point, being a threshold with its two error rates, rather than by a single accuracy figure, because accuracy alone does not say which of the two errors the system is making. On the detection stage, Bochkovskiy, Wang and Liao [3] introduced YOLOv4, which the project names. Their reported figures are given in Table II. It is a general object detector trained on a multi-class dataset; applied to faces it locates them in the frame. Viola and Jones [4] established the classical detection baseline still widely used for this task. For the recognition stage, the modern approach embeds an aligned face into a vector space in which distance corresponds to identity, and Schroff and colleagues [5] give the formulation on which most current systems rest. For presentation attacks, the international standard [6] defines the terminology and the testing methodology for detecting artefacts presented to a biometric sensor.

Table 1. Representative Related Work and its Relation to This Study

Ref.	Year	Focus and scope	Method and validation	Relation to this work
[1]	—	Face recognition for door access	Prototype systems	The application class
[2]	—	Digital lock and home automation	Prototype systems	System context
[3]	2020	YOLOv4 object detector	Benchmark on MS COCO	The detection stage named; Table II
[4]	2001	Rapid object detection	Classical cascade detector	The detection baseline
[5]	2015	Face embedding and verification	Benchmark evaluation	What the recognition stage requires
[6]	—	Presentation attack detection	International standard	Terminology and testing for spoofing

Entries [1] and [2] point to bodies of prototype literature rather than to individual works, as the project record does. Specific works should be named before submission.

1. What the Named Detector Actually Reports

Because the project names YOLOv4 specifically, its published figures are worth stating, both for what they establish and for what they do not. Table 2 gives them.

Table 2. Published YOLOv4 Figures

Quantity	Published value	Conditions
Average precision, MS COCO	43.5 % AP; 65.7 % AP at 50 % overlap	608 pixel input, test-dev 2017
Speed, Volta GPU	62 to 96 frames per second	Tesla V100 or Titan V, 608 down to 416 pixel input
Speed, Pascal GPU	33 to 54 frames per second	Titan X, 1080 Ti or Tesla P100
Speed, Maxwell GPU	23 to 38 frames per second	GTX Titan X or Tesla M40
YOLOv4-tiny variant	22.0 % AP at about 443 frames per second	RTX 2080 Ti

Figures are from the YOLOv4 papers [3]. Three points bear on this project. The average precision is a detection score on a general multi-class benchmark and says nothing about identifying a person. Every speed quoted is

for a discrete graphics processor; the computer used in this project is not specified, and on a processor without a suitable graphics card the frame rate is a small fraction of these values. And the accuracy of a detector places an upper bound on the system only in the sense that an undetected face cannot be recognised; it is not itself a recognition accuracy.

Relations Governing a Biometric Lock

1. Detection Is Not Recognition

The distinction is the foundation of everything that follows and is shown in Fig. 1. A detector answers a question about the image: is there a face, and where. A recogniser answers a question about a person: whose face is this. The second requires the first, but the first does not provide the second.

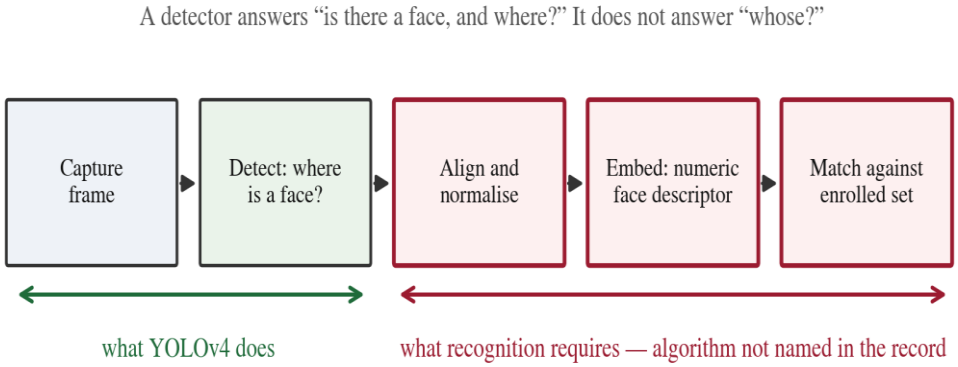

Figure 1: The recognition pipeline. YOLOv4 performs the detection stage; the remaining stages are what identification requires.

Between detection and a decision lie three further stages. The detected face must be aligned and normalised so that pose and scale do not dominate the comparison; it must be converted into an embedding, a numeric descriptor in a space where distance corresponds to identity [5]; and that embedding must be compared against the enrolled set. The record names an algorithm for the first stage and none for the other three, which is the gap identified in Section V-A.

2. The Decision Rule

Given an embedding of the presented face and a set of enrolled embeddings, access is granted when the best match exceeds a threshold,

$$\text{accept if } \max_i s(e, e_i) \geq \tau \qquad (1)$$

Equation (1) is the whole of the security decision, and the threshold in it is the single parameter that sets how the system behaves.

3. The Two Errors

A threshold decision makes two kinds of error,

$$FAR = \Pr\{s \geq \tau \mid \text{impostor}\},$$
$$FRR = \Pr\{s < \tau \mid \text{genuine}\} \qquad (2)$$

A false accept admits a stranger; a false reject locks out a resident. They trade against each other along the threshold, as Fig. 2 shows: tightening the threshold to reduce one raises the other, and no threshold removes both. The point at which they are equal is the equal error rate, which is a convenient single number for comparing systems but is not usually the right operating point for a lock, since the two errors have very unequal consequences.

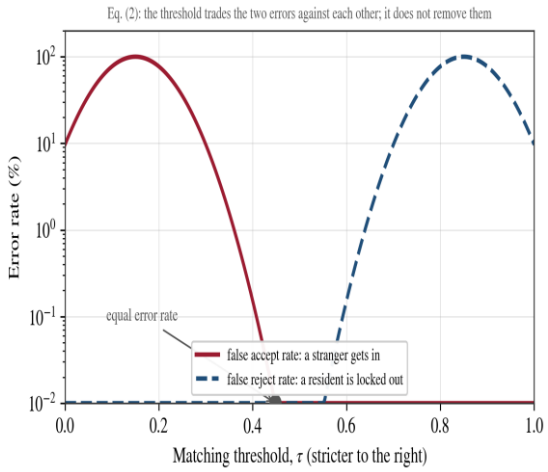

Figure 2: The threshold trade-off of Eq. (2). Curves are illustrative of the shape; they are not measurements of any system.

A single accuracy percentage does not distinguish the two, which is why it is not an adequate report for a lock. A system that never opens has perfect security and zero utility, and would score well on some accuracy definitions.

4. Identification Is Weaker Than Verification

The distinction that matters most for this design is between verification and identification. Verification asks whether a presented face

matches one claimed identity, which is what a phone unlock does after the user has been asserted by possession of the phone. Identification asks which, if any, of the enrolled identities a presented face matches, which is what a door lock must do because a visitor claims nothing.

Under identification, an impostor gets one chance against each enrolled identity, so the system false accept rate accumulates,

$$FAR_{sys} = 1 - (1 - FAR)^N \approx N \cdot FAR \quad (3)$$

Table 3 evaluates this and Fig. 3 plots it. The consequence is that a lock becomes less secure with every additional person enrolled, which is the opposite of what intuition suggests and is not a defect of any particular algorithm.

Table 3. System False Accept Rate Against Number Enrolled (Computed from Eq. (3))

Per-comparison FAR	N = 5	N = 10	N = 20	N = 50
0.01 %	0.05 %	0.10 %	0.20 %	0.50 %
0.1 %	0.50 %	1.00 %	1.98 %	4.88 %
1 %	4.90 %	9.56 %	18.2 %	39.5 %

Computed from Eq. (3), which assumes the comparisons are independent. Real impostor comparisons are correlated, so the true values differ somewhat, but the direction and order of magnitude hold. The per-comparison rates are illustrative operating points, not measurements of any system in this project.

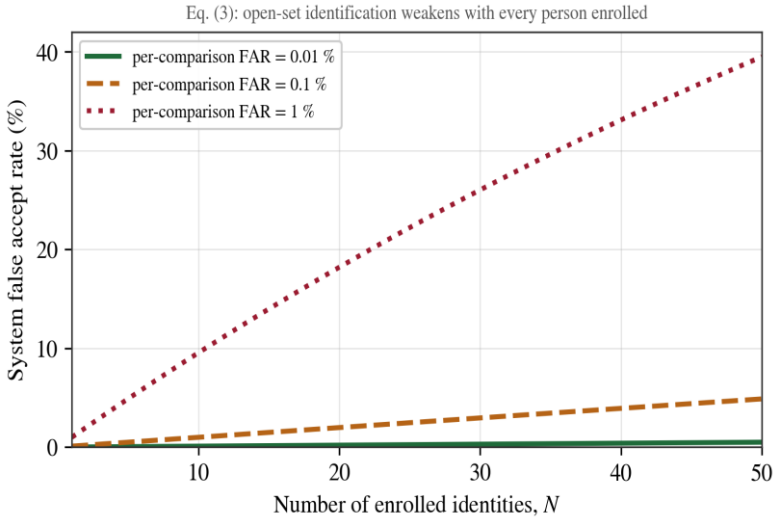

Fig. 3: System false accept rate against the number of enrolled identities, from Eq. (3).

5. What a Rate Means in Practice

A percentage is difficult to act on, so it is worth converting. If impostor faces are presented to the camera at some rate, whether visitors, delivery drivers or passers-by detected at the entrance, the mean interval between false admissions is

$$T_{int} = \frac{1}{r \cdot FAR_{sys}} \quad \text{[days]} \quad (4)$$

Table 4 evaluates this at ten presentations per day and Fig. 4 plots it more generally.

Table 4. Mean Interval Between False Admissions at Ten Impostor Presentations per Day

System false accept rate	Mean interval	Reading
10 %	1 day	A stranger admitted daily
1 %	10 days	Roughly three times a month
0.1 %	100 days	About three times a year
0.05 %	200 days	Under twice a year

Computed from Eq. (4). The presentation rate of ten per day is an assumption chosen to represent a household entrance; the true rate depends on the site and would need to be counted. These are illustrative conversions, not measurements.

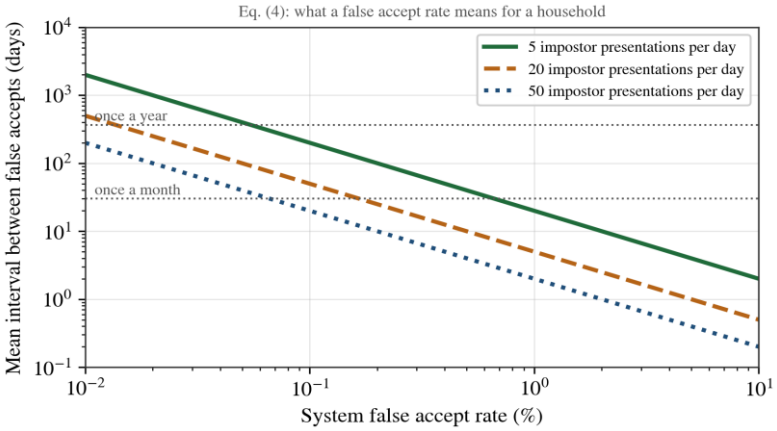

Figure 4: Mean interval between false admissions from Eq. (4), at three impostor presentation rates.

6. Latency

The user-facing performance figure is the time from presenting a face to the lock releasing,

$$T_{e2e} = t_{cap} + t_{det} + t_{emb} + t_{match} + t_{act} \quad (5)$$

comprising capture, detection, embedding, matching and actuation. The detection term is the one the published figures of Table II bear on, and only for the hardware those figures were measured on.

The System

1. Architecture

Fig. 5 shows the system. A webcam at the entrance feeds a computer which runs the detector, holds the enrolled database and applies Eq. (1). On a match, the computer sends a command to the Arduino Uno over a serial link, and the Arduino energises the relay that releases the lock. The dashed line marks the trust boundary discussed in Section VI-B.

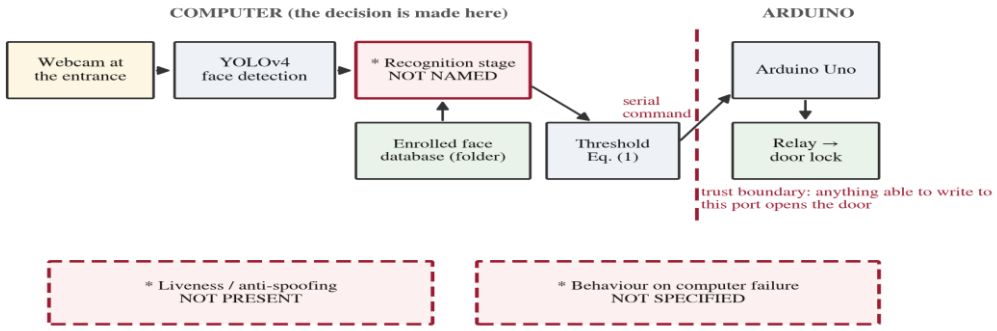

Figure 5: System architecture. The recognition decision is taken on the computer and conveyed to the Arduino as a serial command.

2. Apparatus

Table 5. System as Described, and Status to Be Confirmed

Item or quantity	As described	Status
Camera	Webcam at the entrance	described
Detector	YOLOv4	described
Enrolled database	Face images stored at a path	described
Control unit	Arduino Uno driving a relay	described

Entries marked as not stated or not described refer to the abstract, which is the only source available to this paper. Several may well be addressed in the full report.

What the Lock Claim Requires

1. The Recognition Stage Is Unnamed

The record states that YOLOv4 is used for detecting the faces and obtaining clear images of them, and separately that the camera recognises a face and compares it against the database. The first is a detection task and YOLOv4 performs it. The second is a recognition task, and the algorithm performing it is not named anywhere in the abstract.

This matters because every quantity in Section III belongs to the recognition stage, not the

detection stage. The threshold of Eq. (1), the error rates of Eq. (2) and the accumulation of Eq. (3) are all properties of the matcher. A report that names only the detector has named the part of the system that does not determine its security. Identifying the recognition method is therefore the first thing the full paper must do, and stating its threshold the second.

2. An Accuracy Figure Would Not Be Enough

Should the full record report a recognition accuracy, it is worth stating in advance why that alone would not settle the question. Accuracy conflates the two errors of Eq. (2), and for a lock they are not interchangeable: a false reject is an inconvenience that the resident resolves with a key, while a false accept is the failure the lock exists to prevent. The report should give both rates at one stated threshold, and should state the number of identities enrolled during the trial, because Eq. (3) makes the system rate depend on it.

The scale of that dependence is worth restating from Table III. At a per-comparison false accept rate of 0.1 per cent, which would be a respectable figure for a modest embedding model, a household of ten reaches a system rate of about 1 per cent, and by Table IV that admits a stranger roughly every ten days at a plausible presentation rate. The same algorithm in a verification role, where the user first claims an identity, would be an order of magnitude stronger. The architecture, not the algorithm, produces the weakening.

3. A Photograph Is a Face to a Webcam

A conventional webcam produces a two-dimensional image and cannot distinguish a live face from a printed photograph of it, a phone screen displaying it, or a mask. Detecting such artefacts is a separate function, termed presentation attack detection, and it is the subject of its own standard and testing methodology [6]. No such stage appears in the described design.

The exposure is specific rather than theoretical for this application. A face is a public credential: it is visible to anyone at the door, photographable from the street, and frequently published by its owner on social media. A lock that accepts a photograph is therefore defeated by an attacker who has never met the resident. Adding a liveness stage, whether by challenge-response, by a depth or infrared camera, or by texture analysis, is not an optional refinement for a face-based lock; it is the difference between a credential and a public identifier.

Security, Safety and Availability

Four consequences follow from the architecture and are stated rather than deferred, because a lock is a safety-relevant device.

1. Egress Must Not Depend on the System

The record states that the door is locked unless a face is recognised. That is the correct default for security, and it must not be the arrangement for leaving. Free egress from a dwelling is a life-safety requirement in a fire, and it must not depend on a camera, a computer or a power supply. The inside of the door should open mechanically without any electronic decision, and the record should say so explicitly.

2. The Serial Link Is the Trust Boundary

The recognition decision is made on the computer, and the Arduino receives a command and closes a relay. The Arduino does not verify who sent the command or why. Anything with access to that serial port, whether software on the computer or a device physically substituted for it, opens the door without presenting any face at all. This is a property of the architecture, not a fault in the code, and it means the security of the lock is bounded by the security of the computer. Where that computer is a general-purpose machine used for other things, that bound is a low one.

3. Availability

The recognition runs on a computer, so the residents' ability to enter depends on that computer being powered, booted and running the application. A crash, an update, a power cut or a disconnected webcam leaves the door locked against the people who live behind it. A lock is a device whose failure mode must be designed rather than discovered, and the record does not state what happens in any of these cases.

4. The Enrolled Database Is Personal Data

The system stores images of the residents' faces at a path on a computer. Those images are biometric data about identifiable people, and unlike a password a face cannot be reissued after a compromise. How the database is protected, who can read it, and what happens to it when the system is decommissioned are questions the design should answer. Storing embeddings rather than images, where the recognition method permits it, reduces but does not remove the exposure.

5. Qualitative Position

Table 6 sets out the structural properties that can be asserted, each with the relation or source that supports it.

Table 6. Qualitative Properties and Their Basis

Feature	Property	Basis
YOLOv4 stage	Locates faces; does not identify them	[3], Fig. 1
Threshold decision	Trades false accepts against false rejects	Eqs. (1), (2)
Open-set identification	System error grows with the number enrolled	Eq. (3), Table 3
Two-dimensional camera	Cannot distinguish a face from a photograph of one	[6]
Face as a credential	Public and not reissuable after compromise	Section VI-D
Serial command link	Bounds lock security by computer security	Section VI-B
Computer-hosted decision	Residents' entry depends on the computer running	Section VI-C
Locked-by-default	Correct for entry; must not govern egress	Section VI-A

Limitations

The following limitations bound every claim made in this paper.

- This paper was prepared from the project abstract alone. It makes no statement about what the campaign measured, and the status column of Table V is a set of questions for the authors.
- No false accept rate, false reject rate, accuracy, latency or spoof-resistance figure is reported here, because none was available to this paper.
- The per-comparison error rates in Table III are illustrative operating points chosen to span a plausible range. They are not measurements of any system in this project.
- Equation (3) assumes independent comparisons. Real impostor comparisons are correlated, so the computed system rates are approximations; the direction and order of magnitude hold but the exact values would differ.
- The presentation rate of ten impostors per day in Table IV is an assumption representing a household entrance. The true rate is site-specific and would have to be counted.
- Figure 2 illustrates the shape of the threshold trade-off. The curves are not measurements of any system.
- The YOLOv4 figures in Table II are published benchmark results on general object detection with discrete graphics hardware. They are not face-recognition accuracies and they do not describe the speed of this project's computer, which is unspecified.
- No claim of suitability for fitting to an occupied dwelling is made anywhere in this paper.

Proposed Measurement Framework

Fig. 6 shows the four trials that would produce the figures a lock claim requires. None needs equipment beyond the system itself and a printed photograph.

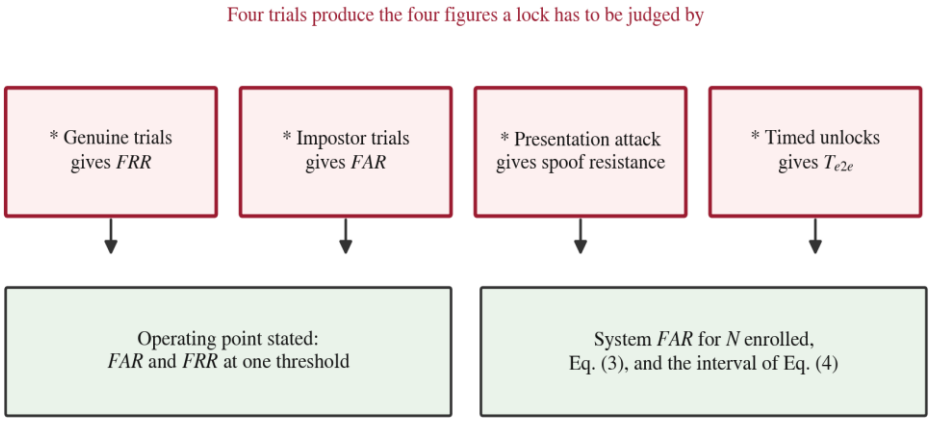

Figure 6: Measurement framework. Four trials produce the four figures a lock has to be judged by.

1. Measurement Plan

Table 7 sets out the trials, and Table 8 records the figures used in this paper with their provenance.

Table 7. Measurement Plan Required Before Any Lock Claim

Trial	Method	What it yields
Name the recognition method	State the alignment, embedding and matching algorithms and the threshold	The parameters of Eq. (1)
Genuine trials	Repeated presentations by each enrolled person across lighting, pose and time of day	FRR at the stated threshold
Impostor trials	Repeated presentations by non-enrolled people	FAR at the same threshold
Report the operating point	Both rates at one threshold, with N stated	A comparable result; enables Eq. (3)
Presentation attack trials	Printed photograph, phone screen, and any other artefact, following the standard methodology [6]	Spoof resistance, or its absence
Timed unlocks	Time from face presented to lock released, with the computer specified	T_{e2e} by Eq. (5)
Failure-mode tests	Power loss, computer crash, camera disconnected, during both entry and egress	Documented and designed failure behaviour
Database handling	State where images or embeddings are held, how protected, and the retention policy	The privacy position of Section VI-D

Table 8. Figures Used in This Paper and Their Provenance

Item	Value	Status
YOLOv4 average precision	43.5 % AP on MS COCO	published [3]
YOLOv4 speed, Volta GPU	62 to 96 frames per second	published [3]
Per-comparison FAR values	0.01, 0.1 and 1 %	illustrative operating points
System FAR at N = 10, FAR = 0.1 %	about 1 %	computed, Eq. (3)
Interval at 1 % and ten presentations a day	about 10 days	computed, Eq. (4)

Published entries are benchmark results for a general object detector. Computed entries follow from the stated relations and illustrative operating points. No measurement of this project's system appears anywhere in this paper.

Conclusion

A face-recognition door-lock system was described, in which a webcam at the entrance feeds a computer running a YOLOv4 detector, enrolled face images are held as a database, and an Arduino Uno drives the relay that releases the lock on a match. The detection, matching and decision relations were consolidated, and the figures such a lock must be judged by were set out.

Three findings follow. First, YOLOv4 is an object detector: it locates faces and does not identify them. Recognition additionally requires alignment, an embedding and a matching step, and the algorithm performing those is not named in the record. Since the threshold, the error rates and their accumulation are all properties of the matcher rather than the detector, the report names the part of the system that does not determine its security.

Second, a lock is characterised by its error rates at a stated threshold, not by an accuracy percentage, and because a door lock performs open-set identification rather than verification the system false accept rate grows with the number enrolled. At a per-comparison rate of 0.1 per cent, a household of ten reaches about 1 per cent and one of fifty about 5 per cent; converted through the presentation rate, a 1 per cent system rate admits a stranger roughly every ten days at ten impostor presentations per day. The weakening comes from the architecture, not from any algorithm, and it is the reason the number enrolled must be reported alongside the error rates.

Third, a two-dimensional webcam cannot distinguish a live face from a printed photograph, and no anti-spoofing stage is described. Since a face is a public credential that cannot be reissued after compromise, liveness detection is not a refinement for a lock of this kind but a

precondition. The paper further identified the serial link to the Arduino as the architecture's trust boundary, noted that the residents' ability to enter depends on a general-purpose computer remaining in service, and stated that egress must never depend on the system. A four-trial measurement plan was given, requiring no equipment beyond the system itself and a printed photograph.

This paper was prepared from the project abstract alone and makes no statement about what the campaign recorded; the status column of Table V is a set of questions for the authors to answer from their full record.

Future Work

Several extensions follow directly from the limitations above.

- Name the recognition method and its threshold-State the alignment, embedding and matching algorithms and the threshold at which access is granted, so that Eq. (1) is fully specified.
- Report an operating point, not an accuracy-Run genuine and impostor trials and report both error rates at one stated threshold, together with the number of identities enrolled during the trial.
- Report the system rate for the enrolled population-Apply Eq. (3) to the measured per-comparison rate and the actual enrolment, and convert it through Eq. (4) so the result is stated in terms a householder can act on.
- Add and test liveness detection-Introduce an anti-spoofing stage and test it against a printed photograph and a phone screen following the standard methodology, reporting the outcome either way.
- Design the failure modes-Decide and document what happens on power loss, computer crash and camera disconnection, for both entry and egress, and confirm that egress is mechanical and independent of the system.
- Harden the trust boundary-Either authenticate the command the Arduino accepts, or state explicitly that the lock's security is bounded by the security of the computer.
- State the database handling-Record where enrolled images or embeddings are stored, how they are protected, who can read them, and what happens to them when the system is decommissioned.
- Measure and report latency-Time the interval from face presented to lock released on the actual hardware, and state that hardware, since the published detector speeds are for discrete graphics processors.

Acknowledgment

The authors thank the Department of Electrical and Electronics Engineering, Stella Mary's College of Engineering, for laboratory facilities and guidance during the project.

References

Automated face recognition systems for door access control, in the prototype and applications literature. A specific work should be named before submission.

Smart digital lock systems for home automation, in the prototype and applications literature. A specific work should be named before submission.

A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "YOLOv4: Optimal speed and accuracy of object detection," arXiv:2004.10934, 2020; and C.-Y. Wang, A. Bochkovskiy, and H.-Y. M. Liao, "Scaled-YOLOv4: Scaling cross stage partial network," arXiv:2011.08036, 2020.

P. Viola and M. Jones, "Rapid object detection using a boosted cascade of simple features," in Proc. IEEE Conf. Computer Vision and Pattern Recognition, 2001.

F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in Proc. IEEE Conf. Computer Vision and Pattern Recognition, 2015, pp. 815-823.

International Organization for Standardization and International Electrotechnical Commission, ISO/IEC 30107, Information technology - Biometric presentation attack detection.