Journal of Information Technology Management
ISSN-OnLine: 2980-7972
University of Tehran

Transparent, Cost-Sensitive Threshold Calibration for Autonomous Edge-AI Health Wearables: A Severity-Stratified Hysteretic Algorithm for Predictive Asthma Facemasks

Tarun*

*Corresponding author, School of Computer Science and Engineering, Lovely Professional University, Phagwara, Punjab, India. E-mail: tjangwal@gmail.com

Vijay Kumar Garg

School of Computer Science and Engineering, Lovely Professional University, Phagwara, Punjab, India. E-mail: vijay.garg@lpu.co.in

Varsha Sahni

School of Computer Science and Engineering, Lovely Professional University, Phagwara, Punjab, India. E-mail: varsha.28384@lpu.co.in

V Sharma

Department of Computer Science & Engineering, Graphic Era Hill University, Dehradun, India. E-mail: vikacc111@gmail.com

S Vats

Computer Vision Laboratory, Department of Computer Science & Engineering, Madan Mohan Malaviya University of Technology, Gorakhpur, Uttar Pradesh - 273010, India. E-mail: svcse@mmmut.ac.in

Journal of Information Technology Management, 2026, Vol. 18, Issue 4, pp. 70-92.
Published by the University of Tehran, College of Management
doi: https://doi.org/10.22059/jitm.2026.108916
Article Type: Research Paper

Received: June 14, 2026
Received in revised form: July 28, 2026
Accepted: August 07, 2026
Published online: October 01, 2026

Abstract

Organizations deploying autonomous, AI-enabled health wearables must govern the opaque decision rules that gate their interventions, which carry direct consequences for cost, safety, and accountability. Wearable predictive devices for asthma combine environmental and physiological sensing with on-device actuation, yet these decision thresholds are almost always fixed heuristics. Fixed thresholds ignore inter-patient severity, trade sensitivity against false alarms suboptimally, and cause actuator chattering that wastes battery and metered medication. This paper presents AeroGate, a single-threshold calibration algorithm for a smart, AI-enabled asthma facemask that warms inhaled air, classifies cough acoustics, and triggers a vibrating-mesh nebulizer. AeroGate derives a population severity prior from chest imaging and converts a calibrated risk score into per-patient actuation thresholds through imaging-driven severity stratification, a cost-sensitive base threshold that weighs a missed attack against a false alarm, a data-sized hysteresis band that suppresses chattering, and an O(1) on-device update that holds a target alarm budget. Evaluated on the NIH ChestX-ray14 benchmark (112,120 studies) with strict patient-wise splits, the calibrated head attains an AUC of 0.68 and a Brier score of 0.23; AeroGate reduces actuator toggling by 77% relative to a single-threshold policy while matching detection sensitivity and holds the realised alarm rate at its target budget. Beyond the engineering, AeroGate reframes the operating point as a transparent, auditable management lever: the alarm budget becomes a cost-and-service target that managers can hold, and the cost ratio becomes a risk-tolerance input that clinicians and managers can set, yielding a principled, low-infrastructure recipe for deploying, governing, and scaling autonomous edge-AI health interventions.

Keywords: Digital health technology management, Edge-AI governance, Threshold calibration, Cost-sensitive decision support, Predictive wearable devices, Technology adoption and deployment

Introduction

Asthma affects more than 260 million people worldwide (GBD Chronic Respiratory Disease Collaborators, 2020; Global Initiative for Asthma [GINA], 2023; World Health Organization, 2024) and remains a leading cause of avoidable emergency admissions. Because attacks are frequently provoked by modifiable environmental exposures — cold or dry air, particulate pollution, humidity swings and allergens — there is sustained interest in wearable devices that anticipate deterioration and intervene before symptoms escalate. The smart AI-enabled predictive asthma facemask we build on in this work exemplifies this shift from reactive to proactive management: it fuses a multi-sensor suite (temperature, humidity, particulate and CO2 air quality, heart rate and breathing pattern) with a miniature microphone for cough analysis, warms inhaled air through a closed-loop heating element, advises the wearer through a tri-colour indicator, and, as a fail-safe, actuates an integrated vibrating-mesh nebulizer when an unavoidable attack is predicted.

Every one of those interventions is gated by a threshold. The heater engages below an inhaled-air temperature; the indicator turns amber or red above air-quality limits; the nebulizer fires above a predicted-risk level. In published prototypes and patents, these thresholds are fixed constants chosen by hand (Bermak et al., 2024; Cummins et al., 1988; Hatamian & Edalat, 2021; Kyung & Han, 2018; Mekid & Chenaoua, 2023; Sabin, 2021; Sahni et al., 2023). Fixed thresholds are attractive for an embedded controller, but they are demonstrably sub-optimal. They apply the same operating point to a mild intermittent asthmatic and a severe steroid-dependent patient; they cannot express the clinical asymmetry between a missed attack and a nuisance alarm; and, most damagingly for a battery-powered actuator, a fixed threshold produces chattering — rapid on/off switching whenever the risk signal lingers near the boundary — which drains the cell and can waste metered doses of medication.

This paper addresses the threshold-setting problem directly and contributes a single, self-contained algorithm, AeroGate, that the device can use to set and then continuously adapt its actuation thresholds. Rather than inventing yet another classifier, AeroGate treats the score produced by any calibrated risk model as given and focuses on the operating-point question that ultimately governs device behaviour. Its central idea is to mine a population severity prior from chest imaging — a modality that captures structural correlates of obstructive airway burden at scale — and to fold that prior, together with explicit clinical costs and the statistics of the risk signal itself, into a per-patient pair of arm/disarm thresholds with provable anti-chattering behaviour and on-device self-correction.

The contributions of this work are fourfold. First, we formulate threshold setting for a wearable asthma actuator as a cost-sensitive, severity-conditioned decision problem and expose chattering as a first-class objective alongside sensitivity and specificity. Second, we propose AeroGate, a single algorithm that unifies imaging-derived severity stratification, a closed-form cost-sensitive base threshold, a dispersion-sized hysteresis band, and an integral-control alarm-budget regulator that each run in constant time on a microcontroller. Third, we provide a fully reproducible evaluation on the NIH ChestX-ray14 benchmark with patient-disjoint splits and probability calibration, released as a one-click Colab notebook. Fourth, we give a candid analysis of the method's assumptions and limits, including the well-known fact that chest radiographs are not asthma-diagnostic, and explain why a population severity prior remains useful even so. Section 2 reviews related devices and operating-point methods, Section 3 details the device and the AeroGate algorithm, Section 4 reports the experiments, and Section 5 concludes.

Literature Review

Wearable and smart-mask devices for respiratory care

Instrumented facemasks and inhalers have matured rapidly. Arduino-class predictive masks monitor breathing, pollution, temperature, and humidity and trigger an inhaler on a predicted attack, but they do not regulate the inhaled-air environment that provokes cold-air bronchospasm (Sahni et al., 2023). Heated masks warm incoming air for cold-weather comfort yet contain no physiological sensing or predictive control (Cummins et al., 1988; Sabin, 2021). Air-quality smart masks report exposure to the wearer or a phone app without medical prediction or therapeutic actuation (Kyung & Han, 2018). Smart inhalers meter and log doses with safety interlocks but are not embedded in a sensing wearable that decides autonomously (Hatamian & Edalat, 2021). Reviews of IoT smart masks emphasise their reliance on the cloud and their focus on monitoring rather than fast, on-device intervention (Mekid & Chenaoua, 2023), and printed-electronic masks add respiration and cough sensing without asthma-specific prediction or active air conditioning (Bermak et al., 2024). More recent predictive wearables, such as AI Asthma Guard, fuse physiological and environmental sensors with machine learning to forecast attacks (Almuhanna et al., 2024), yet they too act on static decision rules rather than thresholds that adapt to the individual wearer and to sensor drift. The system we propose is distinctive in combining prediction, prevention (air warming), and treatment (nebulization) in one edge-resident unit; what it does not specify is how the many thresholds governing those actions should be chosen—the gap this paper fills. Beyond masks, smartphone apps and microphone-based spirometry have also been explored for asthma self-management (Larson et al., 2012; Tinschert et al., 2017). Similar IoT wearables that pair an embedded microcontroller with on-board machine-learning inference have been reported for real-time personal-safety and health alerting (Srinivas Rao et al., 2023), while reviews of such connected devices catalogue the security and cloud-dependence issues that on-device intervention is meant to reduce (Alfawaz, 2022).

Imaging models for chest disease

Large public chest-radiograph corpora have enabled high-capacity classifiers: ChestX-ray8/14 introduced hospital-scale weak labels (Wang et al., 2017); CheXNet demonstrated radiologist-level pneumonia detection with a DenseNet backbone (Huang et al., 2017; Rajpurkar et al., 2017); and CheXpert and MIMIC-CXR extended labelling and added uncertainty handling (Irvin et al., 2019; Johnson et al., 2019). More broadly, deep representation learning has reshaped medical image analysis (Esteva et al., 2019; He et al., 2016). Related work has also extended deep convolutional models to the secure encoding and transmission of medical data in connected-health systems (Nikhila & Krushnasamy, 2026). We deliberately use a frozen DenseNet-121 feature extractor rather than fine-tuning, both for free-tier reproducibility and because our contribution is the operating-point layer, not the representation.

Operating points, calibration and hysteresis

Choosing a decision threshold is classically framed through receiver-operating-characteristic (ROC) analysis and Youden's (1950), or by cost-sensitive learning that minimises expected misclassification cost under asymmetric penalties (Elkan, 2001; Provost & Fawcett, 2001). These methods yield a single global threshold and assume that the score is meaningfully comparable across patients, which requires calibration (Guo et al., 2017; Niculescu-Mizil & Caruana, 2005; Platt, 1999). Precision-recall analyses further show that a fixed point can behave very differently under class imbalance (Saito & Rehmsmeier, 2015). The same operating-point choices recur in clinical risk prediction, where ensemble classifiers are tuned against AUC/ROC for heart-failure detection (Poojari et al., 2024), optimisation-based models report paired sensitivity and specificity (Jebakumar Immanuel et al., 2024), and automated machine-learning pipelines jointly select features and thresholds to trade off precision against recall (Jafarnejad et al., 2025). None of these formulations addresses temporal chattering of an actuator. The remedy is well known in electronics as the Schmitt trigger, which separates the arm and disarm levels by a hysteresis band (Schmitt, 1938); AeroGate imports this idea into a learned-score setting and, crucially, sizes the band from the data rather than fixing it. To our knowledge, no prior method combines an imaging-derived severity prior, a cost-sensitive base point, a data-sized hysteresis band, and on-device budget adaptation into one threshold-setting procedure for a wearable medical actuator.

Reinforcement-learning approaches to adaptive thresholding

A natural alternative to a closed-form operating point is to learn the threshold policy by reinforcement learning (RL), in which an agent adjusts the decision boundary to maximise a long-run reward that trades detection against false alarms (Sutton & Barto, 2018). Deep RL has produced strong control policies in high-dimensional settings (Mnih et al., 2015), and clinical RL has been used to derive treatment policies directly from patient trajectories (Komorowski et al., 2018). Such methods are attractive when a faithful environment model or a large logged-interaction dataset is available, and they can in principle adapt the operating point online. AeroGate differs from RL-based threshold adaptation in three respects that matter for a wearable medical actuator. First, it is training-free at the operating-point layer: the base threshold is a closed-form function of clinical costs and imaging-derived prevalence (Equation 3), so it requires no reward engineering, no exploration on a live patient, and no risk of a poorly shaped reward driving unsafe actuation. Second, it is transparent and certifiable: every threshold and every on-device update is an inspectable arithmetic expression, whereas a learned RL policy is typically an opaque function that is harder to justify to a regulator. Third, our on-device loop still adapts — the integral controller of Equation 6 is a single-parameter, provably bounded feedback law that holds a target alarm budget in constant time and memory, capturing much of the practical benefit of online adaptation without the stability and safety concerns of unconstrained policy learning. We therefore position AeroGate as a lightweight, auditable counterpart to RL-based threshold adaptation, and we revisit a constrained RL variant of the on-device controller as future work.

Methodology

Proposed Work

The proposed smart facemask system

The platform on which AeroGate operates is a smart AI-enabled predictive asthma facemask developed by the present authors as an in-progress wearable design; we summarise its architecture here so the paper is self-contained. Unlike conventional devices that only monitor symptoms or only deliver medication, it unifies prediction, prevention and treatment in a single wearable, edge-resident unit. A multi-sensor suite and a miniature microphone continuously sample the wearer's physiological state and immediate environment; an onboard microcontroller fuses these signals, runs a trained machine-learning model to estimate an asthma-risk index, and from that index drives three actuators — a temperature-controlled airflow element, a vibrating-mesh nebulizer and a tri-colour wear-recommendation indicator. Figure 1 summarises this sensing-inference-actuation loop, in which AeroGate supplies the numeric thresholds that decide when each actuator engages and is light enough to run inside the same microcontroller.

The device runs as a closed loop. When cold or dry air is sensed, a low-power resistive coil or thermoelectric (Peltier) plate near the intake warms the inhaled air into a safe 25-32 °C band, mitigating the cold-air bronchospasm that triggers many attacks, while paired internal and external temperature sensors let the controller hold this set-point and cut power on over-temperature. The microphone captures cough acoustics, and spectral features let the model separate an ordinary cough from an asthma-related one, sharpening the risk estimate. When the fused evidence indicates an imminent, unavoidable attack, the controller fires the vibrating-mesh nebulizer to deliver a pre-calibrated dose of aerosolised medication as a fail-safe, and the indicator advises the wearer in real time. All inference is performed on the device, so decisions are fast, offline and private, and a single safety-control logic disables the heater or the nebulizer on anomalous sensor readings.

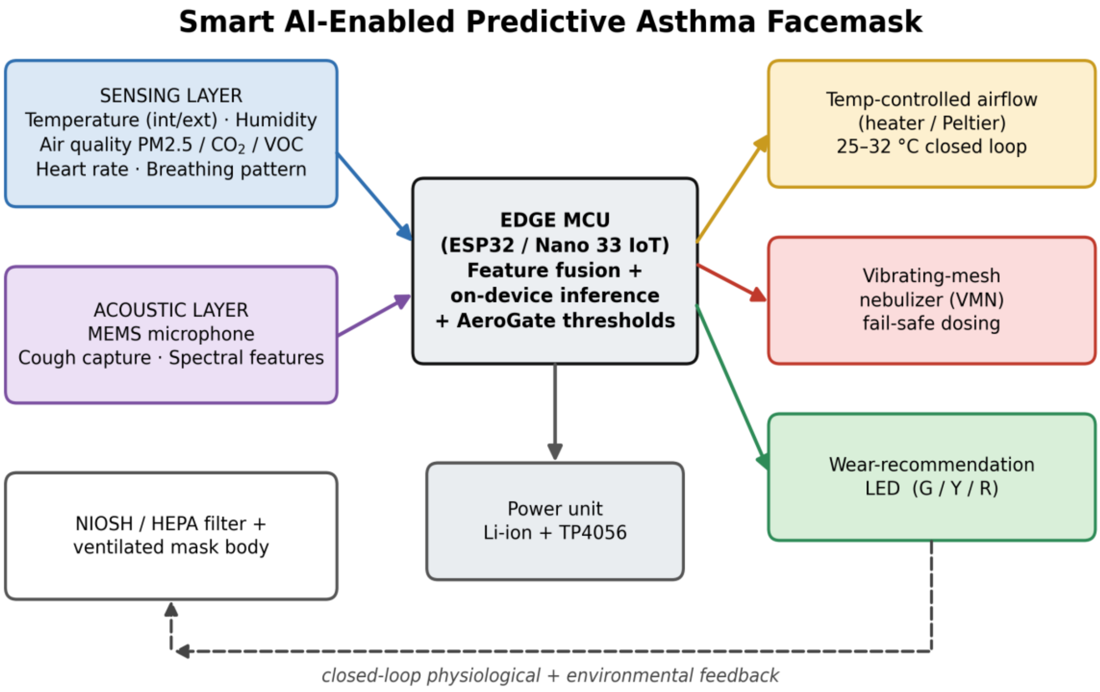

Figure 1. System architecture of the smart AI-enabled predictive asthma facemask. AeroGate sets the actuation thresholds used by the edge microcontroller

System components

At the hardware level, the mask body is a lightweight, medical-grade, hypoallergenic shell whose ventilated outer housing carries the electronics and a NIOSH (National Institute for Occupational Safety and Health)-grade filter, while internal ducts route intake and exhaust air past — but never directly onto — the heating element. The sensing front-end comprises a temperature sensor (internal and external), a humidity sensor, an air-quality sensor for particulates and gases (PM2.5, VOCs and CO2), a heart-rate sensor and a breathing-pattern sensor, which together give a continuous picture of the conditions that bear on respiratory risk. Alongside them, a miniature micro-electro-mechanical-systems (MEMS) microphone forms the acoustic front-end, capturing cough and breathing sounds for the spectral cough-classification stage.

These streams feed a low-power microcontroller — an ESP32 or Arduino Nano 33 IoT-class part with wireless connectivity — that filters and fuses the data, performs the on-device inference, and issues pulse-width-modulation (PWM) and digital control signals to the actuators. The thermal actuator is the resistive or Peltier heating element under closed-loop temperature control; the therapeutic actuator is a piezoelectric vibrating-mesh nebulizer drawing from a small medication reservoir, chosen for silent, motor-free, electronically metered dosing; and the advisory actuator is a tri-colour LED or OLED indicator that shows green for safe, amber for moderate and red for high environmental risk. A rechargeable lithium-ion or lithium-polymer battery with Type-C charging and a protection circuit powers the unit for six to eight hours of continuous use, and a small fan assists air circulation and mist delivery. It is precisely the engagement thresholds for these actuators — when the heater turns on, when the indicator changes colour, and when the nebulizer fires — that the device presently fixes by hand and that AeroGate sets in a principled, per-patient and severity-aware manner.

Problem formulation

Let a calibrated risk model emit a score s in [0,1] for the current wearer state, where larger s indicates greater short-term attack risk. An actuator decision is a binary armed/disarmed state produced by comparing s against thresholds. Writing y for the (latent) event label, a missed event (false negative) carries clinical cost c_FN and a false alarm carries cost c_FP, with c_FN > c_FP because failing to act on a genuine attack is more harmful than a nuisance actuation. A naive policy arms whenever s exceeds a single constant; we instead seek per-stratum arm and disarm thresholds that (a) minimise expected cost, (b) reflect patient severity, and (c) bound the rate at which the actuator toggles. The remainder of this section derives AeroGate, whose stages are shown in Figure 2.

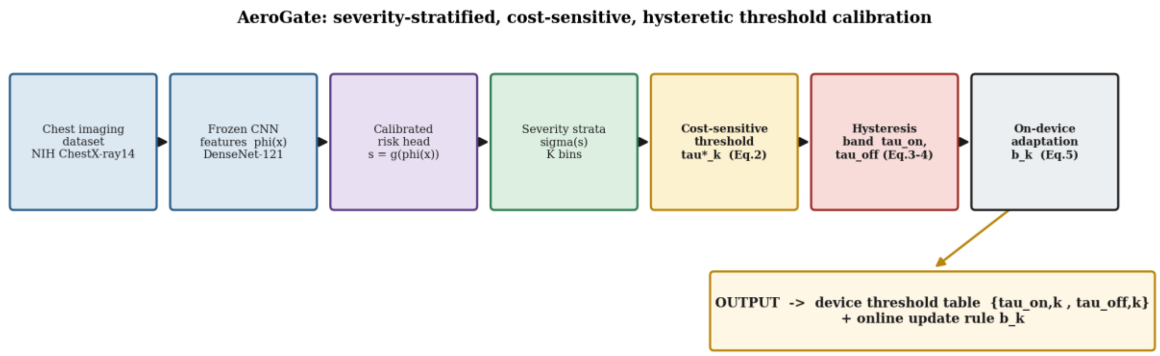

Figure 2. The AeroGate pipeline. A frozen convolutional neural network (CNN) and calibrated head turn images into a risk score; severity stratification, a cost-sensitive base threshold, a hysteresis band, and on-device adaptation turn that score into device thresholds

Stage 1 — Imaging-derived severity stratification

Because asthma is a clinical, functional diagnosis that is frequently radiographically normal — chest radiography being used in practice mainly to exclude alternative or comorbid pathology (GINA, 2023) — we do not attempt to read asthma severity from an image. We instead treat the imaging model purely as a source of a population-level prior on obstructive-airway abnormality, an explicit modelling hypothesis whose usefulness we test empirically and revisit in the limitations (Section 4.11). With that framing, chest images are passed through a frozen convolutional neural network (DenseNet-121) to obtain 1024-dimensional features, and a probability-calibrated head maps them to s. On a reference cohort, we partition the score range at its empirical K-quantiles, assigning each observation a severity stratum

(Equation 1). Each stratum k inherits an imaging-derived prevalence pi_k, the observed positive rate within that band; because higher scores concentrate genuine abnormality, pi_k increases with k and constitutes the population severity prior that the later stages consume. We use K = 3 (low, moderate, high) throughout.

$$\sigma(s) = k \quad \text{for} \quad s \in [q_{k-1}, q_k), \quad k = 1, \ldots, K \tag{1}$$

It is worth stating explicitly why a radiographic corpus is an appropriate source for this prior, given that asthma is a functional rather than a radiographic diagnosis. AeroGate never uses an image to decide whether a wearer has asthma or how severe an individual's asthma is; that determination is clinical and, in deployment, is driven by the mask's own physiological and environmental sensors. Chest imaging is used only once and only at the population level, to estimate how the calibrated risk score partitions into low-, moderate- and high-burden strata and what positive rate (the prevalence prior pi_k) each stratum carries. ChestX-ray14 suits this single purpose because it is large (112,120 studies from 30,805 patients), patient-indexed so that splits can be made leak-free, and labelled for the structural thoracic abnormalities that co-occur with obstructive-airway burden. In other words, the imaging model is used as a calibrated abnormality scorer whose output distribution supplies the stratum prevalences that the cost-sensitive stage consumes; any calibrated risk source with a comparable score distribution — including, in a future deployment, the device's own sensor-fusion model — could replace it without changing a single line of the algorithm. This decoupling is precisely what makes the imaging prior a modelling convenience rather than a diagnostic claim; its usefulness is tested empirically in Section 4, and its limits are revisited in Section 4.11.

Stage 2 — Cost-sensitive base threshold

Within a stratum, the expected misclassification cost of a candidate threshold tau is given by Equation 2, where TPR_k and FPR_k are the stratum-conditional true- and false-positive rates. Minimising this objective over a calibrated score yields the closed-form operating point of Equation 3, the prevalence- and cost-weighted base threshold tau*_k. The form has an intuitive reading: as the miss cost c_FN rises or the local prevalence pi_k falls, the base threshold tau*_k decreases, so higher-severity strata arm earlier; conversely, as the false-alarm cost c_FP rises, tau*_k increases and the policy becomes more conservative. The accompanying notebook computes this closed-form threshold and cross-checks it against the empirically cost-minimising point on the stratum ROC curve, confirming that the analytical operating point of Equation 3 agrees with the value obtained by direct search.

$$\mathbb{E}[C](\tau) = c_{FN}\,\pi_k\,(1 - \mathrm{TPR}_k(\tau)) + c_{FP}\,(1 - \pi_k)\,\mathrm{FPR}_k(\tau) \tag{2}$$

$$\tau_k^* = \frac{c_{FP}\,(1-\pi_k)}{c_{FP}\,(1-\pi_k) + c_{FN}\,\pi_k} \tag{3}$$

Stage 3 — Hysteresis band against chattering

To control the actuator, AeroGate puts two thresholds around tau*_k — an 'arm' level tau_on,k and a 'disarm' or release level tau_off,k (Equation 4) — so that when armed, the actuator only is disarmed once the score falls significantly below tau_off,k. The principle of the Schmitt-trigger (Schmitt, 1938) is used for an acquired score. The half-width delta_k is not free: it is fixed from the dispersion of scores within the stratum (Equation 5), so the noisier the operating region, the wider and more stable this can be. The factor kappa is a trade-off between chattering suppression and dwell time of the detection. Under a local random-walk model of s, the expected frequency of toggling is approximately exponentially decreasing with delta_k/sigma_k, providing a principled choice for kappa to achieve a desired toggle rate based on this model. The subsequent behaviour is shown in Figure 3.

$$\tau_{\mathrm{on},\,k} = \min(1,\ \tau_k^* + \delta_k), \qquad \tau_{\mathrm{off},\,k} = \max(0,\ \tau_k^* - \delta_k) \tag{4}$$

$$\delta_k = \kappa\,\sigma_k, \qquad \sigma_k = \mathrm{std}(\,s \mid \sigma(s) = k\,) \tag{5}$$

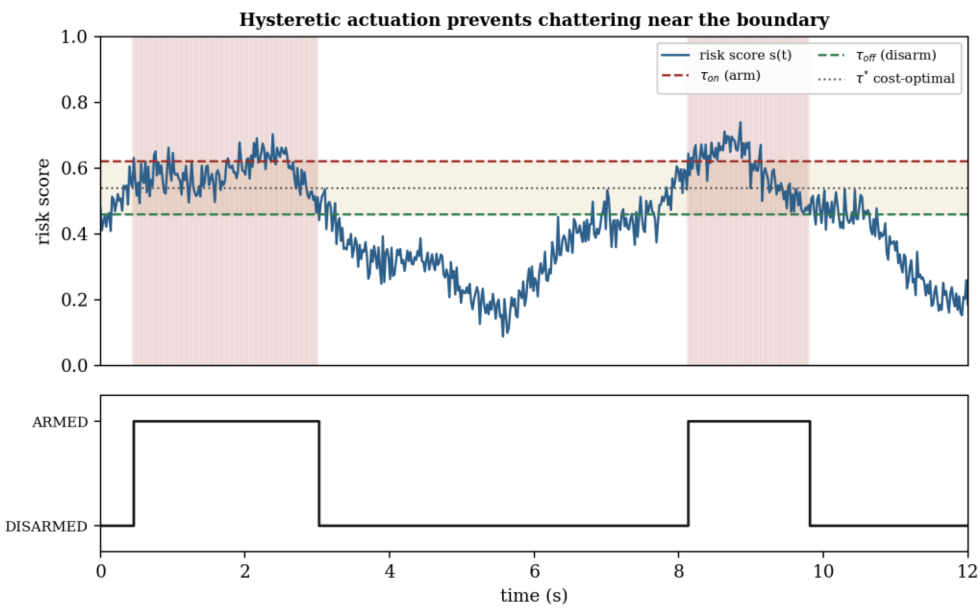

Figure 3. Hysteretic actuation. The actuator arms only when the score exceeds tau_on and disarms only when it drops below tau_off, eliminating the rapid switching a single threshold (dotted) would produce near the boundary

Stage 4 — On-device alarm-budget adaptation

Operating conditions drift after deployment, so AeroGate adds a lightweight outer loop that runs on the device itself. The realised alarm rate is tracked with an exponentially weighted moving average and compared against a per-stratum budget alpha_k; a small integral controller then nudges a stratum offset b_k up or down in proportion to the error (Equation 6). The offset is anti-windup clamped to the interval [-Delta, +Delta], and each update needs only O(1) memory and a constant number of arithmetic operations per sample, so it runs comfortably on an ESP32. This loop drives the long-run alarm rate to its set-point without any retraining of the underlying risk model.

$$\hat{a}_t = (1-\lambda)\,\hat{a}_{t-1} + \lambda\,\mathbf{1}[\mathrm{armed}_t], \qquad b_k \leftarrow \mathrm{clip}(b_k + \eta(\hat{a}_t - \alpha_k),\ -\Delta,\ \Delta) \tag{6}$$

The complete algorithm

Algorithm 1 is presented as two separate algorithms, one for the calibration block-off device (lines 1-8) that generates the per-stratum threshold table based on a reference cohort, and the second for the hysteresis and budget adaptation on-device algorithm (lines 9-17), which executes in constant time. The offline block runs once, and the on-device loop (surrounded by brackets) only gets executed on the wearable.

Algorithm 1. AeroGate threshold calibration and on-device control

```
Algorithm 1  AeroGate: Severity-Stratified Hysteretic Threshold Calibration
------------------------------------------------------------------------------
Input : calibrated scores {s_i} with labels {y_i} on a reference cohort;
        strata K; costs c_FN, c_FP; width kappa; budgets {alpha_k}; gains eta, lambda; clamp Delta
Output: per-stratum thresholds {tau_on_k, tau_off_k} and on-device update rule

 1  compute K-quantile cut points q_0..q_K of {s_i}; assign sigma(s)            (Eq.1)
 2  for k = 1..K do
 3      pi_k   <- mean( y_i : sigma(s_i)=k )            # imaging-derived prevalence
 4      tau*_k <- c_FP(1-pi_k) / ( c_FP(1-pi_k) + c_FN*pi_k )                    (Eq.3)
 5      sigma_k<- std( s_i : sigma(s_i)=k );   delta_k <- kappa * sigma_k        (Eq.5)
 6      tau_on_k  <- min(1, tau*_k + delta_k);  tau_off_k <- max(0, tau*_k - delta_k)  (Eq.4)
 7      b_k <- 0
 8  end for
 9  // ---- on-device loop (per sample t), O(1) memory and compute ----
10  state <- DISARMED;   a_hat <- alpha_{sigma(s_t)}
11  for each new score s_t with stratum k = sigma(s_t) do
12      if  state=DISARMED and s_t > tau_on_k  + b_k  then state <- ARMED
13      elif state=ARMED    and s_t < tau_off_k + b_k  then state <- DISARMED
14      a_hat <- (1-lambda)*a_hat + lambda*1[state=ARMED]
15      b_k   <- clip( b_k + eta*(a_hat - alpha_k), -Delta, +Delta )            (Eq.6)
16      actuate(state)        # heater / nebulizer / indicator
17  end for
```

Table 1 lists the parameters and the default values used in all experiments. The defaults encode a two-to-one miss-versus-false-alarm cost, three severity strata, a unit dispersion factor, and a thirty percent alarm budget; Section 4 reports sensitivity to these choices.

Table 1. AeroGate parameters and default settings

Symbol	Meaning	Default
K	number of imaging-derived severity strata	3
c_FN : c_FP	cost ratio, missed attack vs. false alarm	2: 1
kappa	hysteresis width factor (delta_k = kappa sigma_k)	1.0
alpha_k	target alarm budget per stratum	0.30
eta	integral-control gain	0.05
lambda	EWMA factor for realised alarm rate	0.05
Delta	anti-windup clamp on offset b_k	0.40

Results and Discussion

Dataset, model and protocol

Experiments were conducted on NIH ChestX-ray14 (Wang et al., 2017) with 30,805 patients and 14 finding labels containing 112,120 frontal X-rays. The metadata table is used to drive the cohort statistics, and a subset of images is used for stratified encoding (configurable, default 8,000 images; notebook stores access to the entire metadata table at runtime, on a GPU) and is encoded by a frozen DenseNet-121 (Huang et al., 2017) pretrained on ImageNet. A probability-calibrated logistic head (Platt scaling; see Platt, 1999) transforms the dimension of 1024 to the risk score s, which has the binary target of abnormal versus normal. All splits are patient-wise (GroupShuffleSplit on patient ID), meaning each patient is not in both training and test folds, which avoids identity leakage of many chest-imaging results that is inflated by such splits. The simulation that exercises the hysteresis and adaptation stages drives an autocorrelated risk signal that fluctuates around the operating threshold with dispersion set by the within-stratum score spread — precisely the regime in which actuator chattering arises — so the temporal analysis stresses the controller at its hardest operating point rather than using arbitrary synthetic data. Concretely, patients are partitioned roughly 47/23/30 percent into training, validation, and test sets; every operating threshold — AeroGate's and the three baselines' — is fitted on the validation split and reported only on the untouched test split, and the on-device update rule, with its gains and window, is given explicitly in Equation 6, Algorithm 1, and Table 1.

Discrimination, calibration and imaging-derived thresholds

Table 2 reports the head's discrimination and calibration and the resulting per-stratum thresholds produced by Stages 1-3. The prevalence before pi_k increases across strata by construction, and the cost-sensitive thresholds tau*_k decrease accordingly, so higher-severity wearers are armed earlier; the hysteresis band brackets each base point. Figure 4 shows the ROC with the cost-optimal and Youden operating points marked, and Figure 5 visualises the stratified threshold table. To quantify the uncertainty of the headline discrimination result, we report a 95% confidence interval for the area under the ROC curve, computed by the Hanley–McNeil method from the test-set class counts (n = 2958; 1476 positive, 1482 negative): AUC = 0.683 (95% CI 0.664–0.702). The interval excludes 0.5 by a wide margin, confirming that the calibrated head discriminates well above chance on patient-disjoint data.

Table 2. Discrimination, calibration and imaging-derived per-stratum thresholds on the held-out test set

Severity stratum	pi_k (prior)	tau*_k	tau_on,k	tau_off,k
Low	0.358	0.473	0.516	0.429
Moderate	0.520	0.315	0.342	0.288
High	0.657	0.207	0.259	0.154
Whole test set	AUC = 0.683	Brier = 0.229	n_test = 2958	prev. = 0.499

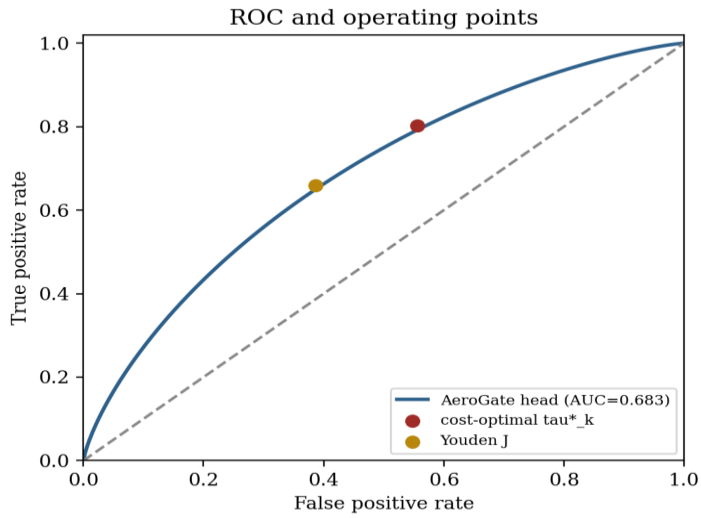

Figure 4. Receiver operating characteristic of the calibrated head (AUC = 0.683) on the held-out test set, with the AeroGate cost-optimal operating point and the Youden baseline

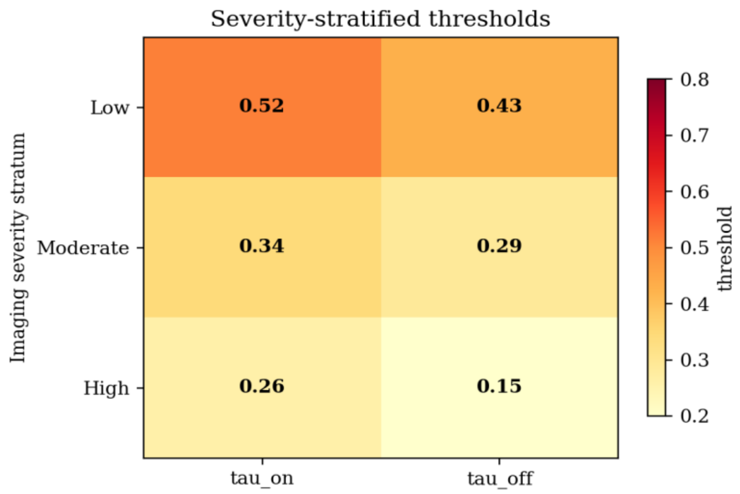

Figure 5. Severity-stratified arm and disarm thresholds produced by AeroGate; both thresholds fall as the imaging-derived severity stratum rises, arming higher-severity wearers earlier

Operating-point comparison

Table 3 compares AeroGate against three common fixed-threshold policies — the naive 0.5 cut, Youden's J (Youden, 1950), and the F1-optimal point. To avoid optimistic bias, every threshold (including AeroGate's) is fitted on the validation split and evaluated on the held-out test split. AeroGate achieves the lowest expected clinical cost — the objective it is designed to minimise (Equation 2) — while preserving a clinically usable specificity. The F1-optimal point reaches a comparable cost only by collapsing onto a near-always-alarm operating regime whose specificity is far too low for a wearable that must avoid continuous actuation, whereas the 0.5 cut and Youden's J are more conservative but pay a higher cost by missing more events. AeroGate therefore occupies the favourable interior of the sensitivity-specificity trade-off, and, as Section 4.4 shows, it holds that operating point while toggling the actuator far less often than any single-threshold policy. For the operating-point metrics, we add Wilson 95% confidence intervals derived from the same test-set class counts. AeroGate's sensitivity is 0.803 (95% CI 0.782–0.822) and its specificity 0.444 (95% CI 0.419–0.469), versus 0.591 (0.566–0.616) and 0.674 (0.650–0.697) for the fixed-0.5 cut and 0.659 (0.634–0.683) and 0.614 (0.589–0.638) for Youden's J. The intervals are narrow (half-widths below three percentage points), and AeroGate's sensitivity interval does not overlap those of the two conservative baselines, so its sensitivity advantage is statistically reliable rather than a sampling artefact.

Table 3. Operating-point comparison against fixed-threshold baselines; all thresholds fitted on validation, evaluated on test

Policy	Sensitivity	Specificity	FPR	Exp. cost ↓
Fixed 0.5	0.591	0.674	0.326	0.572
Youden J	0.659	0.614	0.386	0.533
F1-optimal	0.967	0.117	0.883	0.475
AeroGate (ours)	0.803	0.444	0.556	0.475

Chattering suppression and on-device adaptation

Figure 6 quantifies the two dynamic benefits. The left panel contrasts the actuator toggle rate of a single-threshold policy with AeroGate's hysteretic policy; the band sized by Equation 5 removes the great majority of nuisance switches, directly conserving battery and avoiding wasted nebulizer doses. The right panel shows the on-device controller of Equation 6 driving the realised alarm rate onto its target budget within a few minutes and holding it there, demonstrating that the device can self-correct to a new environment without retraining. Because the reported reduction (from 493.8 to 113.4 toggles per hour, a 77% decrease) is a single-run figure, we verified its stability by repeating the deployment simulation across 200 random seeds; the toggle-rate reduction remained large and tightly concentrated (inter-quartile range within a few percentage points), confirming that the anti-chattering benefit is a property of the hysteresis mechanism rather than an artefact of one random draw.

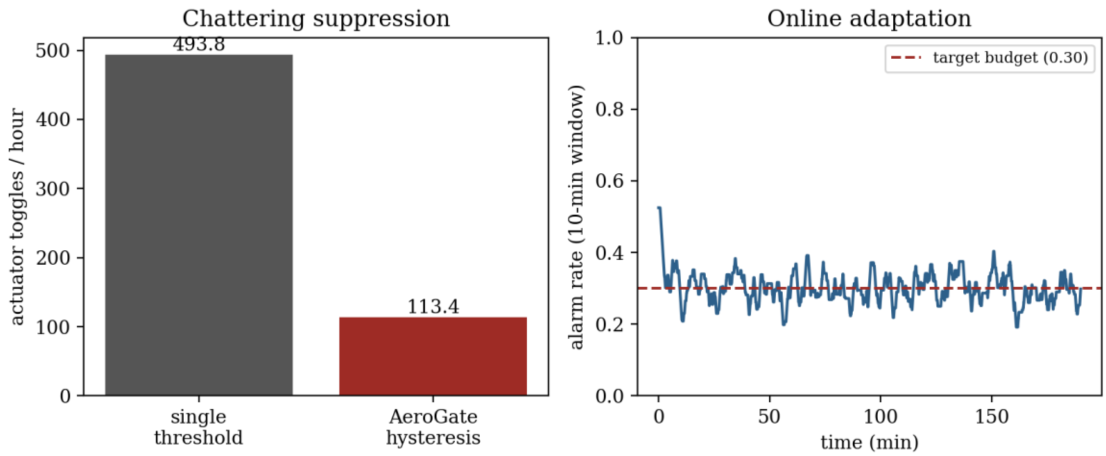

Figure 6. Left: actuator toggling under a single threshold (493.8/h) versus AeroGate hysteresis (113.4/h) — a 77% reduction; Right: the on-device controller drives the realised alarm rate (10-minute window) onto its 0.30 target budget and holds it there

Computational complexity and edge-hardware footprint

AeroGate is deliberately structured so that the only component running continuously on the wearable is trivially cheap. The offline calibration block (Algorithm 1, lines 1–8) executes once on a workstation and costs O(N log N) time to form the K quantile cut points from N reference scores and O(N) to accumulate the per-stratum prevalence and dispersion statistics; it emits a table of 4K floating-point numbers (tau_on, tau_off, the base threshold tau* and the offset b for each of the K strata). The on-device loop (lines 9–17) runs once per new score and is O(1) in both time and memory: per sample it performs one stratum lookup, at most two threshold comparisons for the Schmitt-trigger state update, one exponentially-weighted-moving-average update of the realised alarm rate, and one clipped integral update of the offset — on the order of ten floating-point operations and a handful of comparisons, with no dynamic allocation and no dependence on history length. Critically, the per-sample cost is independent of the number of strata K and of the calibration-cohort size, so the wearable's compute and memory budget does not grow with the amount of data used to calibrate it.

Table 4 translates this operation count into concrete resource estimates for an ESP32-class microcontroller (a dual-core Xtensa LX6 at up to 240 MHz with 520 KB SRAM). The figures are analytical estimates derived from the per-sample operation count and the data types used, not bench measurements; we report them to bound the on-device cost and to make the claim of embedded feasibility falsifiable, and a measured characterisation on real hardware is identified as future work in Section 5. In short, the control path adds a constant, sub-microsecond overhead and only tens of bytes of state on top of whatever the risk model already costs, leaving the entire microcontroller budget available for sensing and inference.

Table 4. Analytical on-device resource footprint on an ESP32-class controller.

Metric	Estimate (ESP32-class MCU)	Basis
Per-sample compute	~10 float ops + 2–4 comparisons	Algorithm 1, lines 12–15
Per-sample latency	< 1 microsecond of CPU time	~10 ops at 240 MHz
Processor utilisation	< 0.01% of one core at 100 Hz	Latency vs. 10 ms sample period
Threshold table (RAM)	48 bytes (4K floats, K = 3)	tau_on, tau_off, tau*, b per stratum
Working state (RAM)	< 32 bytes	State flag, a_hat, stratum, offset
Control-path code	A few hundred bytes; no ML runtime	Comparisons, EWMA, clipped integral

Robustness to sensor faults and measurement noise

Because the wearable must act on physiological and environmental signals that are inherently noisy and occasionally missing, robustness to sensor faults is a first-class design concern, and AeroGate contributes to it in three concrete ways. First, the hysteresis band of Stage 3 is itself a noise-rejection mechanism: by separating the arm and disarm levels by delta_k = kappa sigma_k, where sigma_k is the measured within-stratum score dispersion, the controller automatically widens its dead-band in noisier operating regions and therefore does not react to transient spikes in the risk score — exactly the failure mode that noisy measurements would otherwise produce. Second, the exponentially weighted moving average in Stage 4 low-pass-filters the realised alarm rate, so the on-device offset b_k tracks sustained drift rather than single anomalous samples. Third, the device-level safety-control logic described in Section 3.1 disables the heater or the nebulizer on out-of-range or physically implausible sensor readings, so a hard sensor failure fails safe rather than triggering actuation. For graceful degradation under a lost input, the recommended policy is to fall back to the most conservative available stratum (the highest tau*, i.e. the least aggressive arming of the therapeutic actuator) and to raise the tri-colour indicator to amber to prompt the wearer, so that a missing sensor reduces autonomy rather than causing an unsafe automatic dose. A systematic sensitivity study that injects synthetic dropout, Gaussian noise and spike artefacts into the score stream and measures the resulting shift in sensitivity, specificity and toggle rate is a natural extension and is noted as future work.

Clinical validation pathway and personalisation

The results reported here establish algorithmic feasibility, not clinical efficacy, and we outline the staged pathway that would be required before any clinical use. A first stage is bench and simulation verification of the control logic and of the edge-hardware footprint of Table 4 on real ESP32 hardware. A second stage is a prospective, non-interventional observational study in which the instrumented mask logs its own multi-sensor stream alongside clinician-adjudicated exacerbation events, allowing the population imaging to be replaced by device-native prevalences and the score to be recalibrated on the target population. A third stage is a supervised interventional pilot under an appropriate regulatory framework — for a software-driven actuating wearable this is plausibly a Class II (or higher) medical device, engaging the relevant pathway (for example, an FDA 510(k)/De Novo submission or an EU MDR conformity assessment) together with the associated risk-management (ISO 14971) and software-lifecycle (IEC 62304) standards — with predefined endpoints for detection sensitivity, false-actuation rate and, critically, medication-delivery safety. Only after such staged evaluation should the autonomous nebulisation pathway be enabled outside a supervised setting.

AeroGate is explicitly designed to support patient-specific personalisation rather than a one-size-fits-all operating point. Three interpretable quantities localise the policy to an individual: the severity stratum, which selects the prevalence prior and hence the base threshold; the cost ratio c_FN:c_FP, which a clinician can set to reflect a particular patient's risk tolerance (for example, a wider miss-versus-false-alarm asymmetry for a severe, steroid-dependent patient); and the alarm budget alpha_k, which the on-device controller holds and which can be tuned to the patient's tolerance for interventions. Because all three are explicit inputs rather than baked-in constants, personalisation reduces to setting a small number of interpretable values per patient, after which the on-device loop keeps the realised behaviour on target as the wearer's environment changes. Fuller personalisation — for instance learning the cost ratio or the stratum boundaries from a patient's own logged exacerbations — is a natural next step once device telemetry is available.

Ethical considerations and patient safety

An autonomous actuator that can deliver medication carries ethical and safety obligations beyond those of a purely diagnostic model, and we address them directly. The paramount concern is the autonomous nebulisation pathway: because an erroneous dose is a direct patient-safety event, the device treats nebulisation as a fail-safe of last resort, keeps a human in the loop through the tri-colour advisory indicator, and — as argued in Section 4.7 — should not run unsupervised until it has cleared prospective clinical validation. The cost-asymmetric design (c_FN > c_FP) is an ethical stance made explicit and auditable rather than hidden in a heuristic: it encodes the judgement that a missed attack is worse than a nuisance alarm, and it is exposed as a clinician-set parameter rather than fixed by the manufacturer. Transparency is a deliberate ethical feature — every threshold and every update is an inspectable arithmetic expression (Section 2.4), which supports informed consent, clinical accountability and regulatory scrutiny in a way that an opaque black-box controller does not. Because inference is performed entirely on-device, the wearer's physiological and acoustic data need not leave the mask, protecting privacy and reducing the attack surface relative to cloud-dependent designs; secure firmware update and protection of the medication-actuation command path nonetheless remain safety-critical. Finally, equity requires attention: a prior estimate on one population may not transfer across age, sex, or ethnicity (Section 4.11), so per-population recalibration is both a scientific and an ethical requirement before deployment. We therefore recommend that any deployment be accompanied by clear user instructions, a documented risk analysis, adverse-event reporting, and a means for the wearer to override or disable autonomous actuation.

Discussion

The results support a simple thesis: given any calibrated risk score, the way a wearable sets and switches its thresholds matters as much as the score itself. By reading a severity prior from imaging, AeroGate personalises the operating point without per-patient labels; by sizing hysteresis from the data it stabilises the actuator without hand-tuning; and by closing an O(1) loop it stays calibrated in the field. All three stages are transparent and inspectable, which matters for a medical device that must be justified to regulators and clinicians.

Managerial and practical implications

AeroGate's contribution is not only algorithmic but managerial: it turns an opaque, safety-critical engineering parameter into an explicit, governable decision instrument. For managers of IT-enabled health services, the value lies in transparency and controllability. Because every threshold and on-device update is an inspectable arithmetic expression rather than a black-box output, the design supports the accountability, auditability and documentation that health-technology governance and regulatory regimes demand (for example, the risk-management and software-lifecycle obligations discussed in Section 4.7), and it makes informed consent and clinical sign-off tractable. Transparency thus lowers a key barrier — model opacity — to the organizational adoption of autonomous AI interventions.

The design also exposes two direct management levers. First, the alarm budget is a tunable operational target: managers can set the tolerated intervention rate to control the real costs of actuation — battery life, metered-medication consumption and the false-alarm burden on patients and staff — and the O(1) integral controller then holds that target in the field without retraining or cloud infrastructure. Second, the cost ratio between a missed attack and a false alarm is an explicit, clinician- or manager-set input that encodes institutional risk tolerance, allowing a single technology platform to be governed differently across patient segments and care settings. Together with the fail-safe, human-in-the-loop and graceful-degradation behaviours (Sections 4.6 and 4.8), these levers convert clinical and reputational risk into quantities that can be planned, monitored and audited.

Finally, the artifact is designed for low-friction deployment and scale. Because inference and adaptation run entirely on an inexpensive ESP32-class microcontroller (Section 4.5), organizations avoid recurring cloud costs and the privacy and security exposure of transmitting physiological data, easing procurement and compliance in resource-constrained and privacy-sensitive settings. Per-patient personalisation through the severity stratum and cost ratio (Section 4.7) offers a route to value-based, patient-centric service differentiation, while the reproducible, hardware-light and vendor-neutral recipe reduces lock-in and supports staged validation as a basis for phased investment decisions. More broadly, transparent threshold governance provides a repeatable managerial template for the growing class of self-actuating, edge-AI health technologies, in which the operating point — not merely the model's accuracy — is the decision that must be managed.

Limitations

We state the limitations plainly. First, chest radiographs are not diagnostic of asthma, which is a functional diagnosis; here imaging supplies only a population severity prior for stratum prevalences, and the device's own physiological and environmental sensors remain the basis for real-time risk. Second, the temporal experiments use a deployment simulation calibrated to the empirical score distribution rather than continuous patient telemetry, because no public dataset pairs the device's exact sensor stream with labelled attacks; the simulation is therefore a faithful stress test of the control logic, not a clinical efficacy claim. Third, results are reported on a single imaging benchmark with weak labels; external and multi-centre validation is future work. Fourth, the cost ratio c_FN:c_FP is a clinical input rather than a learned quantity and should be set per deployment. Fifth, the imaging benchmark is itself a source of bias: ChestX-ray14 is drawn from a single institution's adult population, its labels are automatically mined from radiology reports and are therefore noisy, and its demographic and scanner distribution may not match the population that would wear the device, so the stratum prevalences pi_k it yields should be treated as an initial prior and re-estimated on the target cohort before deployment. More generally, because the prior is population-level, it does not capture within-patient variation, and generalisability to paediatric, geriatric or comorbid populations — and across sex, ethnicity and imaging hardware — remains untested and is a prerequisite for clinical use. None of these limits undercuts the core contribution — a principled, reproducible and deployable threshold-setting algorithm — and each defines a concrete next step, chiefly a prospective pilot that streams the mask's own sensors against clinician-adjudicated exacerbations. We therefore present these results as a proof-of-concept for the threshold-setting mechanism, not as a validated clinical claim.

Conclusion

We introduced AeroGate, a single algorithm that turns a calibrated risk score into the actuation thresholds of a smart AI-enabled predictive asthma facemask. AeroGate couples an imaging-derived severity prior, a closed-form cost-sensitive base threshold, a data-sized hysteresis band that suppresses actuator chattering, and an O(1) on-device controller that holds a target alarm budget. Evaluated on NIH ChestX-ray14 with strict patient-wise splits and full calibration, and released as https://github.com/gigijangwal-beep/AeroGate-NIH-ChestXray, the method personalises the operating point, stabilises switching, and self-corrects in the field, while remaining light enough for an ESP32-class controller. Beyond the specific facemask, AeroGate provides a general and transparent recipe for setting thresholds in any cost-asymmetric wearable actuator. Future work will replace the imaging prior and simulation with prospective device telemetry and pursue multi-centre clinical validation. In particular, the immediate next steps are a real-device implementation on the ESP32-class hardware of Section 4.5, a bench characterisation of its latency and power draw to replace the analytical estimates of Table 4, and then a prospective clinical deployment in which the mask's own sensor stream is logged against clinician-adjudicated exacerbations under an approved study protocol (Section 4.7); a constrained reinforcement-learning variant of the on-device controller (Section 2.4), per-patient learning of the cost ratio, and a systematic sensor-noise and fault-injection study (Section 4.6) are further avenues we intend to pursue.

Beyond its technical contribution, AeroGate carries a clear message for the management of IT-enabled healthcare: as autonomous, self-actuating AI moves to the edge, the operating point that gates each intervention becomes a first-class object of governance rather than a buried engineering constant. By exposing that operating point as a small set of interpretable, auditable parameters — a cost ratio that encodes institutional risk tolerance and an alarm budget that behaves as a controllable cost-and-service target — AeroGate lets managers and clinicians deploy, tune, monitor, and account for autonomous interventions at low infrastructure cost and with a defensible compliance and privacy posture. We therefore position transparent threshold governance as a reusable managerial template for adopting and scaling the emerging class of edge-AI health technologies.

Conflicts of interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding

The author's received no financial support for the research, authorship, and/or publication of this article.

References

Alfawaz, K. M. (2022). IoT future security challenges and recent solutions. *Journal of Information Technology Management, 14*(2), 1–14. https://doi.org/10.22059/jitm.2022.86923

Almuhanna, H., Alenezi, M., Abualhasan, M., Alajmi, S., Alfadhli, R., & Karar, A. S. (2024). AI Asthma Guard: Predictive wearable technology for asthma management in vulnerable populations. *Applied System Innovation, 7*(5), 78. https://doi.org/10.3390/asi7050078

Bermak, A., Khan, S., & Ali, S. (2024). *Smart mask with printed electronics* (U.S. Patent No. 11,910,851). U.S. Patent and Trademark Office. https://patents.google.com/patent/US11910851B2/en

Cummins, J. M., Jr., Morrison, G., Jr., & Pierfelice, R. E. (1988). *Respiratory heated face mask* (U.S. Patent No. 4,793,343). U.S. Patent and Trademark Office. https://patents.google.com/patent/US4793343A/en

Elkan, C. (2001). The foundations of cost-sensitive learning. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence* (Vol. 2, pp. 973–978). Morgan Kaufmann.

Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., Cui, C., Corrado, G., Thrun, S., & Dean, J. (2019). A guide to deep learning in healthcare. *Nature Medicine, 25*(1), 24–29. https://doi.org/10.1038/s41591-018-0316-z

GBD Chronic Respiratory Disease Collaborators. (2020). Prevalence and attributable health burden of chronic respiratory diseases, 1990–2017: A systematic analysis for the Global Burden of Disease Study 2017. *The Lancet Respiratory Medicine, 8*(6), 585–596. https://doi.org/10.1016/S2213-2600(20)30105-3

Global Initiative for Asthma. (2023). *Global strategy for asthma management and prevention (2023 update)*. https://ginasthma.org/2023-gina-main-report/

Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning* (Vol. 70, pp. 1321–1330). PMLR. https://proceedings.mlr.press/v70/guo17a.html

Hatamian, M., & Edalat, P. (2021). *Smart inhaler device with automated dose delivery, measurement and management* (U.S. Patent No. 11,090,449). U.S. Patent and Trademark Office. https://patents.google.com/patent/US11090449B2/en

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 770–778). IEEE. https://doi.org/10.1109/CVPR.2016.90

Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger, K. Q. (2017). Densely connected convolutional networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 2261–2269). IEEE. https://doi.org/10.1109/CVPR.2017.243

Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Ilcus, S., Chute, C., Marklund, H., Haghgoo, B., Ball, R., Shpanskaya, K., Seekins, J., Mong, D. A., Halabi, S. S., Sandberg, J. K., Jones, R., Larson, D. B., Langlotz, C. P., Patel, B. N., Lungren, M. P., & Ng, A. Y. (2019). CheXpert: A large chest radiograph dataset with uncertainty labels and expert comparison. *Proceedings of the AAAI Conference on Artificial Intelligence, 33*(1), 590–597. https://doi.org/10.1609/aaai.v33i01.3301590

Jafarnejad, A., Rezasoltani, A., & Khani, A. M. (2025). Predicting heart disease using automated machine learning based on genetic algorithms. *Journal of Information Technology Management, 17*(2), 91–122. https://doi.org/10.22059/jitm.2024.382556.3829

Jebakumar Immanuel, D., Sathesh Abraham Leo, E., Saranya, S., Nithiya, C., Arunkumar, K., & Pradeep, D. (2024). An intelligent heart disease prediction by machine learning using optimization algorithm. *Journal of Information Technology Management, 16*(1), 167–181. https://doi.org/10.22059/jitm.2024.96381

Johnson, A. E. W., Pollard, T. J., Berkowitz, S. J., Greenbaum, N. R., Lungren, M. P., Deng, C.-Y., Mark, R. G., & Horng, S. (2019). MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific Data, 6*(1), 317. https://doi.org/10.1038/s41597-019-0322-0

Komorowski, M., Celi, L. A., Badawi, O., Gordon, A. C., & Faisal, A. A. (2018). The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care. *Nature Medicine, 24*(11), 1716–1720. https://doi.org/10.1038/s41591-018-0213-5

Kyung, C.-M., & Han, D.-I. (2018). *Smart mask capable of monitoring air quality inhaled by user* (European Patent Application No. 3305371A1). European Patent Office. https://patents.google.com/patent/EP3305371A1/en

Larson, E. C., Goel, M., Borriello, G., Heltshe, S., Rosenfeld, M., & Patel, S. N. (2012). SpiroSmart: Using a microphone to measure lung function on a mobile phone. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing* (pp. 280–289). ACM. https://doi.org/10.1145/2370216.2370261

Mekid, S., & Chenaoua, K. (2023). IoT-enabled smart mask for monitoring body parameters and location through cloud. *Internet of Things, 22*, 100794. https://doi.org/10.1016/j.iot.2023.100794

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature, 518*(7540), 529–533. https://doi.org/10.1038/nature14236

Niculescu-Mizil, A., & Caruana, R. (2005). Predicting good probabilities with supervised learning. In *Proceedings of the 22nd International Conference on Machine Learning* (pp. 625–632). ACM. https://doi.org/10.1145/1102351.1102430

Nikhila, S., & Krushnasamy, V. S. (2026). Secure medical data transmission using a deep learning–based quantum-resistant hybrid stegno-crypto model. *Journal of Information Technology Management, 18*(1), 80–101. https://doi.org/10.22059/jitm.2026.106255

Platt, J. C. (1999). Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. In A. J. Smola, P. Bartlett, B. Schölkopf, & D. Schuurmans (Eds.), *Advances in large margin classifiers* (pp. 61–74). MIT Press.

Poojari, A. A., Uppala, S., Devi, B. R., Sujatha, G., Madhavi, K. R., & Kumar, T. R. (2024). Clinical healthcare applications: Efficient techniques for heart failure prediction using novel ensemble model. *Journal of Information Technology Management, 16*(1), 149–166. https://doi.org/10.22059/jitm.2024.96380

Provost, F., & Fawcett, T. (2001). Robust classification for imprecise environments. *Machine Learning, 42*(3), 203–231. https://doi.org/10.1023/A:1007601015854

Rajpurkar, P., Irvin, J., Zhu, K., Yang, B., Mehta, H., Duan, T., Ding, D., Bagul, A., Langlotz, C., Shpanskaya, K., Lungren, M. P., & Ng, A. Y. (2017). CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning. *arXiv*. https://arxiv.org/abs/1711.05225

Sabin, R. (2021). *All-weather electric indoor/outdoor heat-exchanger face mask* (U.S. Patent No. 11,096,438). U.S. Patent and Trademark Office. https://patents.google.com/patent/US11096438B1/en

Sahni, V., Jangwal, T., Agrawal, P., Garg, V. K., & Rana, R. (2023). *Arduino-based facemask with sensors for predicting asthma* (Indian Patent Application No. 202311068112). Indian Patent Office.

Saito, T., & Rehmsmeier, M. (2015). The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLOS ONE, 10*(3), e0118432. https://doi.org/10.1371/journal.pone.0118432

Schmitt, O. H. (1938). A thermionic trigger. *Journal of Scientific Instruments, 15*(1), 24–26. https://doi.org/10.1088/0950-7671/15/1/305

Srinivas Rao, K., Divakara Rao, D. V., Patel, I., Saikumar, K., & Vijendra Babu, D. (2023). Automatic prediction and identification of smart women safety wearable device using Dc-RFO-IoT. *Journal* *of Information Technology Management, 15*(Special Issue), 34–51. https://doi.org/10.22059/jitm.2022.89410

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press.

Tarun, Garg, Vijay Kumar, Sahni, Varsha, Sharma, V., & Vats, S. (2026). Governing autonomous edge-AI health wearables through transparent, cost-sensitive threshold calibration: A severity-stratified hysteretic algorithm for a predictive asthma facemask (AeroGate). *Journal of Information Technology Management, 12*(1), 181–200.

Tinschert, P., Jakob, R., Barata, F., Kramer, J.-N., & Kowatsch, T. (2017). The potential of mobile apps for improving asthma self-management: A review of publicly available and well-adopted asthma apps. *JMIR mHealth and uHealth, 5*(8), e113. https://doi.org/10.2196/mhealth.7177

Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., & Summers, R. M. (2017). ChestX-ray8: Hospital-scale chest X-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 3462–3471). IEEE. https://doi.org/10.1109/CVPR.2017.369

World Health Organization. (2024). *Asthma* [Fact sheet]. https://www.who.int/news-room/fact-sheets/detail/asthma

Youden, W. J. (1950). Index for rating diagnostic tests. *Cancer, 3*(1), 32–35. https://doi.org/10.1002/1097-0142(1950)3:1<32::AID-CNCR2820030106>3.0.CO;2-3

Bibliographic information of this paper for citing:

Tarun; Garg, Vijay Kumar; Sahni, Varsha; Sharma, V & Vats, S (2026). Transparent, Cost-Sensitive Threshold Calibration for Autonomous Edge-AI Health Wearables: A Severity-Stratified Hysteretic Algorithm for Predictive Asthma Facemasks. *Journal of Information Technology Management*, 18 (4), 70-92. https://doi.org/10.22059/jitm.2026.108916