# GoosebumpsEdge: Wearable On-Device Skin Imaging for Piloerection Detection


Hiroki Ota
Nara Institute of Science and Technology
Japan
tesula22@gmail.com

Alexander Marquardt
Nara Institute of Science and Technology
Japan

Abdallah El Ali
Centrum Wiskunde & Informatica
Utrecht University
The Netherlands

Felix Dollack
Nara Institute of Science and Technology
Japan

Kiyoshi Kiyokawa
Nara Institute of Science and Technology
Japan

Bernhard E. Riecke
School of Interactive Arts + Technology (SIAT), Simon Fraser University
Canada

Monica Perusquia-Hernandez
Nara Institute of Science and Technology
Japan
perusquia@ieee.org


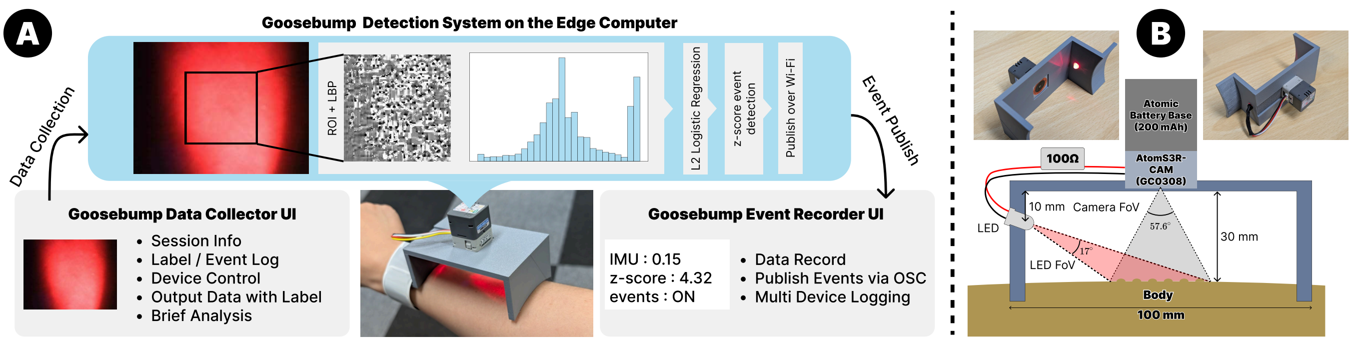


**Figure 1: (A) System overview of GoosebumpsEdge. A wearable camera captures skin-surface images when the skin is lit with a red LED, extracts a fixed Region of Interest (ROI), computes lightweight texture features, estimates goosebump probability on-device, and publishes events to external applications via OSC through a PC Hub. The Collector supports real-time interval labeling and session-wise data export. (B) GoosebumpsEdge: Wearable skin imaging hardware concept that stabilizes distance and illumination using off-the-shelf components plus 3D-printable parts.**


## Abstract

Goosebumps, or piloerection, provide a visible physiological cue that can reflect affective and thermoregulatory responses. We present GoosebumpsEdge, a reproducible ecosystem for wearable skin-surface imaging and lightweight on-device goosebumps detection. Our system consists of (1) a stable imaging hardware jig assembled from four off-the-shelf parts and 3D-printable components, (2) an AtomS3R-CAM firmware for robust image capture with optional Inertial Measurement Unit (IMU) logging, (3) a real-time labeling Collector, and (4) an experiment-facing Hub that publishes goosebumps events via Open Sound Control (OSC). Pilot data from three participants (10 sessions total) indicates that goosebump-related signals are reliably reflected in simple texture features, including FFT-based and Laplacian-based measures. Our approach reliably distinguished target patterns (goosebump events) from background data. Using a Local Binary Pattern (LBP) histogram with L2-regularized logistic regression, we achieved an ROC-AUC of 0.84 and an average precision of 0.65 on a held-out session, with a median ROC-AUC of 0.74 across labeled sessions. The on-device pipeline runs in real time at approximately 5 FPS and converts per-frame probabilities into discrete events using EMA smoothing, hysteresis, and optional z-score-based session adaptation. Together, results demonstrate the feasibility of real-time, on-device goosebumps detection using




*AHs 2026, Okinawa, Japan*



ACM ISBN 979-8-4007-2351-3/2026/03
https://doi.org/10.1145/3795011.3797377

wearable imaging. By releasing all hardware designs and software components as open source, we aim to support further exploration of piloerection as an embodied physiological signal in augmented human systems.

## Keywords

Goosebumps, piloerection, sensing, wearable, edge devices, skin imaging

**ACM Reference Format:**
Hiroki Ota, Alexander Marquardt, Abdallah El Ali, Felix Dollack, Kiyoshi Kiyokawa, Bernhard E. Riecke, and Monica Perusquia-Hernandez. 2026. GoosebumpsEdge: Wearable On-Device Skin Imaging for Piloerection Detection. In *The Augmented Humans International Conference 2026 (AHs 2026), March 16–19, 2026, Okinawa, Japan.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3795011.3797377

## 1 Introduction

Goosebumps, or piloerection, are a visually observable skin response that can accompany emotional, cognitive, and thermoregulatory processes [4, 7, 11]. For most individuals, piloerection is not under direct voluntary control, although voluntary induction has been reported [4]. In augmented human and VR research, measuring goosebumps provides an objective, time-resolved marker that complements self-reports (e.g., presence, affect, comfort) and behavioral measures (e.g., task performance, interaction logs) and is less susceptible to demand characteristics and deliberate response strategies than self-report alone. This is useful because self-report is intermittent and retrospective [3], whereas goosebumps can be continuously and non-invasively captured, enabling moment-by-moment analysis and closed-loop adaptation of stimuli. However, previous solutions typically stream raw images to a host PC for processing, which increases system footprint and complicates deployment. Accordingly, on-device inference enables a fully wireless configuration while improving privacy and scalability by transmitting only detection results rather than raw images. However, many existing approaches rely on specialized sensing setups or offline processing, which limits portability and reuse across studies [12]. Therefore, we aim to make goosebump sensing practical by enabling wearable skin-surface imaging and real-time inference on a low-cost edge device (Figure 1 (A)). This work extends the GooseLab project [1] by emphasizing on-device inference, maintainable implementation, and an experiment-friendly interface via OSC.

**Contributions.** First, we present GoosebumpsEdge, a reproducible end-to-end system for wearable skin-surface imaging and on-device goosebumps detection. The system includes a stabilized imaging module, robust AtomS3R-CAM firmware for reliable capture and device control with optional IMU logging, a real-time interval-labeling Collector that exports session-wise data, and a Hub that publishes discrete events via OSC to interactive applications and experiments. Second, we prove the concept that goosebump-related signals can be captured by simple texture cues (an FFT-derived index and the mean absolute Laplacian response), and demonstrate a lightweight edge pipeline using LBP histograms and L2-regularized logistic regression with on-device, z-score-based event detection.

[1]http://www.goosecam.de/gooselab.html

## 2 Related Work

Objective and continuous piloerection measurement has been studied using optical recording and frequency-domain analysis, demonstrating that goosebumps can be quantified from skin appearance [1]. Image-based approaches further analyze skin texture changes via feature extraction, enabling non-contact emotion-related measurements [12]. Alternative wearable approaches include flexible capacitive sensors that detect skin deformation directly [5], although these require skin contact and are limited to a single measurement site. Piloerection detection has applications in affective computing, including aesthetic experiences like musical chills [9], and has been investigated in Virtual Reality (VR) contexts linking awe experiences to goosebumps [10], though empirical evidence suggests the need for objective, continuous measurement beyond self-report [6]. Prior work typically relies on webcams on a host PC for goosebump detection, whereas our approach runs inference on an edge microcontroller, enabling a fully wireless configuration.

## 3 Implementation

**Hardware Implementation.** Goosebumps manifest as subtle, millimeter-scale texture changes, requiring stable imaging with consistent scale, focus, and illumination. Our wearable jig maintains a fixed camera-to-skin distance and reduces viewpoint variability, while controlled illumination stabilizes image statistics across frames. The current build uses four commercially available components: an AtomS3R-CAM (GC0308), an ATOMIC Battery Base (200 mAh), a 5 mm red LED (625 nm, OSHR5111A), and a 100 Ω resistor, along with a 3D-printable enclosure. To reduce ambient light leakage, black felt is wrapped around the measurement region during recording. Figure 1 (B) illustrates the hardware.

The Atom firmware captures RGB565 frames, extracts a fixed center ROI, and supports runtime control, status monitoring, and debugging.To improve robustness, camera power-cycling and initialization retries are used to mitigate occasional sensor probe failures at boot, and can optionally log synchronized IMU data.

Because continuously computing FFTs on a microcontroller is costly, we do not adopt frequency-domain measures commonly used in prior work [1, 2]. Instead, we use LBP [8] as an inexpensive texture descriptor based on integer comparisons over a small local neighborhood, and estimate goosebumps probability with an L2-regularized logistic regression model: $p = \sigma(\mathbf{w}^\top \mathbf{x} + b)$. To obtain discrete events, per-frame probabilities are smoothed using an Exponential Moving Average (EMA) with hysteresis. We compute an on-device z-score using online EMA estimates of the mean and variance to reduce session-dependent score shifts. These statistics are updated recursively from the start of each session without storing an explicit sliding window, and thus incorporate the entire past with exponential forgetting.

**Software Implementation.** The Collector supports real-time interval labeling by recording timestamped onset/offset events and exports session-wise data (frames, metadata, and event logs) for reproducible analysis. The Hub registers devices and interacts with three lightweight HTTP endpoints: `/status` for reporting device state (e.g., streaming/recording flags, FPS, battery/WiFi status, and last error), `/control` for updating runtime settings (e.g., start/stop

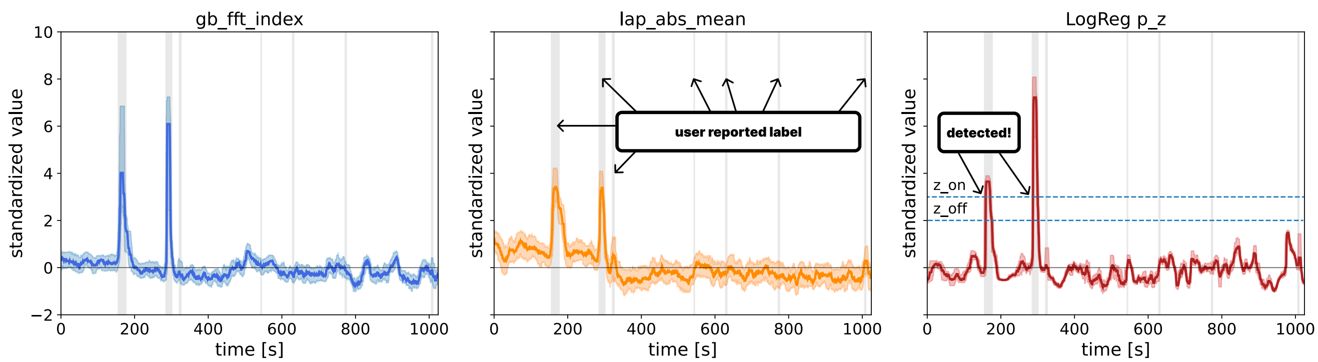


**Figure 2: Example time series from P2 (additional session) recorded during horror movie viewing. Each panel shows signals standardized to zero mean and unit variance, with an interquartile range (IQR) band and median trace over a sliding window. Left: FFT-derived texture index. Center: Laplacian absolute mean. Right: z-score of the LBP-based logistic regression output used for event triggering. The gray-shaded regions indicate the intervals that the user labeled as "goosebumps."**

streaming or recording, LED/illumination, and capture parameters), and `/snapshot` for capturing a single still frame of the skin-surface ROI on demand.The Hub periodically polls `/status`, triggers `/snapshot` when requested by the UI, and publishes discrete goosebumps on/off events via OSC to applications such as Unity. Our implementation and scripts are publicly available at https://github.com/hatodove22/Goosebumps_edge.git.

## 4 Proof of concept

To demonstrate the system's goosebump detection, three participants (P1–P3; 2 male, 1 female; age $M = 37.3$, $SD = 1.53$) watched self-selected videos to maximize the likelihood of eliciting goosebumps. Each participant completed three recording sessions (nine sessions total). In addition, we collected one extra session from P2 and reserved it *a priori* as a fully held-out test session (10 sessions total). This held-out session was excluded from all model development (including cross-validation, feature definition, training, and any model/threshold selection) and was used only for final, external evaluation and for visualizing a representative time series (Fig. 2).

Across the nine non-held-out sessions (19270 frames), the total recording duration was 1389 s (median 160 s/session), of which 194 s (14%) were labeled as goosebumps by the user. Goosebumps frequency varied across participants and sessions, ranging from 0–6 events (median 2), reflecting both individual differences and session-specific recording length. Frames and event logs were stored per session without mixing, enabling session-wise analysis and debugging.

We first verified that goosebump-related signals are captured by simple texture cues, including an FFT-derived texture index and the mean absolute Laplacian response. We then assessed generalization in two complementary ways. First, to estimate typical performance and quantify between-session variability without overfitting, we performed leave-one-session-out cross-validation across the nine non-held-out sessions, where all frames from one session were held out for testing while the remaining sessions were used for training. Second, after fixing the pipeline based only on these nine sessions, we evaluated all methods on P2's held-out session as an external test set never used for training or model selection. FFT/Laplacian baselines and the *LBP + LogReg* model scores were evaluated on the same held-out frames. On P2's held-out test session, the texture baselines showed frame-level separability with ROC-AUC values of 0.78 (FFT-derived index) and 0.809 (mean absolute Laplacian response). We then trained an *LBP + LogReg* model that matched the firmware feature definition and obtained an ROC-AUC of 0.84 and an average precision of 0.65 on the same held-out session. Over the nine labeled sessions used in cross-validation, the median ROC-AUC was 0.74, indicating variability across sessions and individuals. We confirmed real-time inference on AtomS3R-CAM at approximately 5 FPS under practical operation. Figure 2 shows a representative time series from the held-out test session (P2; additional session) that aligns texture cues and the LogReg z-score with labeled intervals.

## 5 Discussion and Future Work

Our results show the feasibility of acquiring and identifying skin-surface images that contain goosebump information with GoosebumpsEdge, and that lightweight inference can run on-device in real time. Compared to prior image-based approaches that rely on host PCs and offline analysis, our implementation shifts piloerection sensing toward event-based, privacy-preserving deployment that can be embedded directly into interactive systems. Still, eliciting goosebump events reliably for an extensive evaluation remains challenging [4, 7, 11]. Additional experiments are needed to evaluate generalization across people and environments where skin appearance, lighting, and motion artifacts can differ substantially. In this setting, accurately estimating event boundaries is more challenging than frame-level discrimination because boundary timing depends on post-processing choices, threshold calibration, and labeling uncertainty. While our results indicate that simple, interpretable texture cues are sufficient for detecting goosebump-related signals, improving robustness at the event level remains an important direction. To this end, IMU logs offer a promising low-cost extension to detect motion segments and suppress intervals likely affected by motion artifacts, without increasing on-device computational complexity. As a next step, we plan to integrate the sensing module into a head-mounted display configuration to support VR experiments with synchronized, real-time measurements of physiological events, enabling closed-loop interaction.

## Acknowledgments

M.P-H. was supported by JSPS KAKENHI 25K21250. H.O. was supported by JST SPRING, Japan Grant Number JPMJSP2140.

## References

[1] Mathias Benedek and Christian Kaernbach. 2010. Objective and continuous measurement of piloerection. *Psychophysiology* 47, 5 (2010), 989–993. doi:10.1111/j.1469-8986.2010.01003.x

[2] Mathias Benedek and Christian Kaernbach. 2011. Physiological correlates and emotional specificity of human piloerection. *Biological psychology* 86, 3 (2011), 320–329.

[3] Abdallah El Ali, Monica Perusquía-Hernández, Pete Denman, Yomna Abdelrahman, Mariam Hassib, Alexander Meschtscherjakov, Denzil Ferreira, and Niels Henze. 2020. MEEC: First Workshop on Momentary Emotion Elicitation and Capture. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–8.

[4] James A. J. Heathers, Kirill Fayn, Paul J. Silvia, Niko Tiliopoulos, and Matthew S. Goodwin. 2018. The voluntary control of piloerection. *PeerJ* 6 (2018), e5292. doi:10.7717/peerj.5292

[5] Jaemin Kim, Dae Geon Seo, and Young-Ho Cho. 2014. A flexible skin piloerection monitoring sensor. *Applied Physics Letters* 104, 25 (2014).

[6] Jonathon McPhetres and Andrew Shtulman. 2021. Piloerection is not a reliable physiological correlate of awe. *International Journal of Psychophysiology* 159 (2021), 88–93.

[7] Jonathon McPhetres and Janis H. Zickfeld. 2022. The physiological study of emotional piloerection: A systematic review and guide for future research. *International Journal of Psychophysiology* 179 (2022), 6–20. doi:10.1016/j.ijpsycho.2022.06.010

[8] Timo Ojala, Matti Pietikäinen, and Topi Mäenpää. 2002. Multiresolution Gray-Scale and Rotation Invariant Texture Classification with Local Binary Patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24, 7 (2002), 971–987. doi:10.1109/TPAMI.2002.1017623

[9] Claire Pelofi, Michal Goldstein, Dana Bevilacqua, Michael McPhee, Ellie Abrams, and Pablo Ripollés. 2021. Chiller: A computer human interface for the live labeling of emotional responses. In *NIME 2021*. PubPub.

[10] Denise Quesnel and Bernhard E. Riecke. 2018. Are You Awed Yet? How Virtual Reality Gives Us Awe and Goose Bumps. *Frontiers in Psychology* 9 (2018), 2158. doi:10.3389/fpsyg.2018.02158

[11] Etain A. Tansey and Christopher D. Johnson. 2015. Recent advances in thermoregulation. *Advances in Physiology Education* 39, 3 (2015), 139–148. doi:10.1152/advan.00126.2014

[12] Mihiro Uchida, Rina Akaho, Keiko Ogawa-Ochiai, and Norimichi Tsumura. 2019. Image-based measurement of changes to skin texture using piloerection for emotion estimation. *Artificial Life and Robotics* 24, 1 (2019), 12–18. doi:10.1007/s10015-018-0435-0

accepted 6 February 2026