



OPEN ACCESS

EDITED BY

Zhuo Zou,
Fudan University, China

REVIEWED BY

Jie Yang,
Westlake University, China
George Brayshaw,
University of Bristol, United Kingdom

*CORRESPONDENCE

Mohammadreza Koolani
✉ Koolani.mohammad@gmail.com

RECEIVED 29 October 2025

REVISED 12 December 2025

ACCEPTED 24 December 2025

PUBLISHED 23 January 2026

CITATION

Koolani M, Bamford S, Trunin P,
Müller-Cleve SF, Lo Preti M, Mastrogiovanni F,
Beccai L and Bartolozzi C (2026) An
event-based opto-tactile skin.
Front. Neurosci. 19:1735068.
doi: 10.3389/fnins.2025.1735068

COPYRIGHT

© 2026 Koolani, Bamford, Trunin,
Müller-Cleve, Lo Preti, Mastrogiovanni, Beccai
and Bartolozzi. This is an open-access article
distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The
use, distribution or reproduction in other
forums is permitted, provided the original
author(s) and the copyright owner(s) are
credited and that the original publication in
this journal is cited, in accordance with
accepted academic practice. No use,
distribution or reproduction is permitted
which does not comply with these terms.

An event-based opto-tactile skin

Mohammadreza Koolani^{1*}, Simeon Bamford¹, Petr Trunin^{2,3},
Simon F. Müller-Cleve¹, Matteo Lo Preti⁴,
Fulvio Mastrogiovanni⁵, Lucia Beccai² and Chiara Bartolozzi¹

¹Event-Driven Perception for Robotics, Istituto Italiano di Tecnologia, Genoa, Italy, ²Soft BioRobotics and Perception Lab, Istituto Italiano di Tecnologia, Genoa, Italy, ³The Open University Affiliated Research Centre at Istituto Italiano di Tecnologia (ARC@IIT), Istituto Italiano di Tecnologia, Genova, Italy, ⁴Advanced Robotics Centre, Department of Mechanical Engineering, National University of Singapore, Singapore, Singapore, ⁵TheEngineRoom, Department of Informatics, Bioengineering, Robotics and Systems Engineering, University of Genoa, Genova, Italy

This paper presents a neuromorphic, event-driven tactile sensing system for soft, large-area skin, based on the Dynamic Vision Sensors (DVS) integrated with a flexible silicone optical waveguide skin. Instead of repetitively scanning embedded photoreceivers, this design uses a stereo vision setup comprising two DVS cameras looking sideways through the skin. Such a design produces events as changes in brightness are detected, and estimates press positions on the 2D skin surface through triangulation, utilizing Density-Based Spatial Clustering of Applications with Noise (DBSCAN) to find the center of mass of contact events resulting from pressing actions. The system is evaluated over a 4,620 mm probed area of the skin using a meander raster scan. Across 95 % of the presses visible to both cameras, the press localization achieved a Root-Mean-Squared Error (RMSE) of 4.66 mm. The results highlight the potential of this approach for wide-area flexible and responsive tactile sensors in soft robotics and interactive environments. Moreover, we examined how the system performs when the amount of event data is strongly reduced. Using stochastic down-sampling, the event stream was reduced to 1/1,024 of its original size. Under this extreme reduction, the average localization error increased only slightly (from 4.66 mm to 9.33 mm), and the system still produced valid press localizations for 85 % of the trials. This reduction in pass rate is expected, as some presses no longer produce enough events to form a reliable cluster for triangulation. These results show that the sensing approach remains functional even with very sparse event data, which is promising for reducing power consumption and computational load in future implementations. The system exhibits a detection latency distribution with a characteristic width of 31 ms.

KEYWORDS

dynamic vision sensors, event-based sensing, neuromorphic engineering, optical skin, soft robotics, stereo vision, tactile sensing

1 Introduction

Soft tactile sensors detect and interpret mechanical stimuli by leveraging compliant materials' inherent flexibility and adaptability. In the past few years, various transduction mechanisms have been extensively explored, including piezoresistive, capacitive, inductive, and optoelectronic approaches (Bartolozzi et al., 2016). Among these, optical sensing stands out due to its wide sensitivity range, excellent reliability, and inherent resistance to electromagnetic interference (Lo Preti et al., 2022a). Unlike many tactile sensors relying on dense wiring or embedded electronics within the sensing area, non-array soft optical waveguides benefit from a sensitive area devoid of wires and rigid parts, delivering fully

flexible and adaptable sensing surfaces (Kamiyama et al., 2004; Li et al., 2022). This feature mainly benefits tactile systems requiring conformity to complex or changing surfaces (Chorley et al., 2009; Guo et al., 2022). Moreover, spatial data in optical soft sensors can be retrieved from broad regions without embedded electronics, relying on data processing, which supports seamless integration into robot platforms and versatile applications in environments that are dynamic, unstructured, or subject to frequent change (Guo et al., 2022; Faris et al., 2023).

Two main types of optical tactile sensors can be identified, namely Frustrated Total Internal Reflection (FTIR) and vision-based systems. On the one hand, FTIR sensors operate by detecting changes in light propagation within an optical waveguide caused by deformation or contact, thus allowing a precise measurement of pressure and shape through photodetection (Trunin et al., 2025; Preti et al., 2025). Vision-based optical sensors, on the other hand, utilize cameras to capture deformation or changes in the sensor surface (Shimonomura, 2019). Vision-based optical sensors, which capture deformation or surface changes via cameras, can be divided into frame-based and event-based approaches. Frame-based vision tactile sensors typically capture images at fixed rates to monitor surface deformation or marker displacements, enabling spatially detailed force and shape sensing. For example, Kamiyama et al. (2004) use color-coded markers embedded in a transparent elastic layer to map three-dimensional contact forces, aiding delicate robot tasks. The compact DIGIT sensor (Lambeta et al., 2020) implements low-cost, high-resolution tactile feedback with a modular elastomer design suitable for robust in-hand manipulation and industrial use. Alternative optical methods, such as the retrographic sensor by Johnson and Adelson (2009) leverage photometric stereo to detect surface texture and shape, while Trueeb et al. (2020) employ multiple cameras and adopts Deep Learning (DL) for 3D force reconstruction over larger areas, enabling scalable robot skins. Inspired by human fingertip anatomy, the TacTip sensor (Chorley et al., 2009) uses internal marker displacement within an artificial papillae structure to achieve sensitive edge detection and manipulation without any embedded electronics.

Dynamic Vision Sensors (DVS)—also known as event cameras—are inspired by biological vision systems and detect brightness changes asynchronously, providing high-speed, low-latency, and data-efficient sensing (Lichtsteiner et al., 2008; Gallego et al., 2022). Event-based approaches have driven advances in tactile perception for robots. Event-based approaches have driven advances in data-efficient tactile perception for robots. Recent reviews highlight that neuromorphic tactile sensors emulate biological mechanoreceptors using spike-based, event-driven encoding, producing sparse, low-latency, and energy-efficient touch signals Liu et al. (2025); Kang et al. (2024). For example, Faris et al. (2023) integrate DVS cameras into a soft robotic finger, achieving proprioception and slip detection via real-time event-based heat maps with a 2 ms latency, enhanced by DL techniques and a flexible fin-ray finger design. Similarly, Naeini et al. (2020) demonstrates a successful application of DVS sensors for high-sensitivity force estimation and material classification, highlighting reduced computational and energy demands compared to frame-based methods in dynamic scenarios.

The hybrid DAVIS sensor (Brandli et al., 2014), which generates as output both asynchronous events and synchronous frames, is employed by Rigi et al. (2018) with silicone substrates to detect slip and vibration, therefore showcasing suitability for rapid, high-resolution tactile feedback. Building on marker-based tactile sensing, NeuroTac combines the TacTip design (Lepora, 2021) with DVS sensors to enable object recognition, shear detection, grasp stabilization, and manipulation for anthropomorphic robots Ward-Cherrier et al. (2020). In a complementary approach, Sferrazza and D'Andrea (2019) developed a high-resolution optical sensor using marker tracking and DL for real-time force prediction on soft surfaces. More recently, Funk et al. (2024) introduced Evetac, an event-based optical tactile sensor leveraging 1,000 h event readout rate to sense vibrations up to 48 h, reconstruct shear forces, drastically reduce data rates versus RGB sensors, and demonstrate data-driven slip detection within a robust closed-loop grasp controller. In a complementary approach, Sankar et al. (2025) demonstrate a biomimetic prosthetic hand with three layers of neuromorphic tactile sensors, enabling compliant grasping and achieving $\approx 99.7\%$ accuracy in texture discrimination.

The approaches mentioned above focus on small-scale sensors for fingertips, with limited examples of large-area tactile skins. Additionally, many of these solutions rely on high-resolution vision sensors, which often require substantial computational resources and the use of advanced data-driven techniques, such as Deep Neural Networks (DNNs) or Reinforcement Learning (RL), to process the vast amount of generated data. While promising, these approaches may face scalability, energy efficiency, and real-time processing challenges, especially for large-scale, adaptive soft robot systems. As an extreme example of multimodal touch sensing, Lambeta et al. (2024) built an artificial fingertip with approximately 8.3 million taxels capable of capturing fine pressures, high-frequency vibrations (up to 10 kh), thermal cues, and chemical traces, all processed by on-chip AI.

To the best of our knowledge, we present the first system integrating stereo DVS sensors with an optical skin to perform contact localization via triangulation. Unlike prior works mostly emphasizing force estimation on small fingertip areas, our approach provides a lightweight solution for the localization of contacts, scalable to large flexible skin areas, using minimal computation without relying on either DNNs or complex marker tracking. However, contact-point localization has also been studied acoustically, for example, Lee et al. (2024) introduced “SonicBoom,” in which a distributed microphone array mounted on a robot link pinpoints single taps with sub-centimeter accuracy even under occlusions.

This study builds on Lo Preti et al. (2022b)'s Multitouch Soft Optical Waveguide Skin (MSOWS), a marker-free, optical skin that achieves large-area tactile sensing without embedded electronics in the sensing area. The design combines graded-stiffness Polydimethylsiloxane (PDMS) layers and a Near-infrared (NIR) optical system to sense pressure, providing adaptability to complex shapes and durability against repeated mechanical stress. MSOWS uses pulsed light sources and repetitive sequential scanning through photodiodes integrated around the skin perimeter. The rate of this scan introduces inherent latency into the detection process, which can only be reduced by increasing the scan rate and therefore the

volume of generated data. In this work, we leverage the low-latency change detection capabilities of event-driven photoreceptors to detect the changes in illumination due to tactile contacts on the silicone surface. As a first step to substitute the array of photodiodes in Lo Preti et al. (2022b), we conducted a preliminary study using a pair of DVS integrated in the silicone layer for contact localization.

As a potential skin covering, applicable to larger areas, this setup allows for an efficient triangulation and spatial analysis of touch positions, therefore contributing to the development of scalable and adaptive tactile sensors in soft robots. Unlike the above-mentioned approaches, which rely on DNNs or heavy data-driven algorithms, this work uses a simple, lightweight triangulation for localizing pressure locations on the silicone surface, which is crucial for real-time applications.

An overview of the performance metrics of State of the Art (SOTA) optical tactile sensors is summarized in Table 1, comparing with our approach and showing their hardware, the measurements, the algorithms, and the task-relative metrics.

2 Materials and methods

2.1 Sensor design and manufacturing process

This paper introduces a new design for a vision-based sensor using two DVS cameras, partly inspired by classic soft optical sensors integrating LEDs and receivers within soft, transparent silicone substrates. This sensor features a square 100 mm per side and 4 mm thick silicone layer, embedding two DVS cameras (Prophesee EVK1 VGA Gen3.1 sensor, resolution: $u = 640$, $v = 480$), positioned at the corners of one edge. Furthermore, 32 NIR LEDs are mounted along the remaining three edges, all pointing to the center of the edge with the two cameras. This setup helps enhance the visibility of the pressure distribution applied to the silicone layer. The LEDs emit light into the elastic silicone element, where deformation within the material affects the light intensity, which is then detected by the DVS cameras. Using two cameras allows for stereopsis, therefore enabling a 2D localization of pressure locations on the skin.

The sensor design is shown in Figure 1. The base (1), printed from Acrylonitrile Butadiene Styrene (ABS) on an Ultimaker S3 (Ultimaker, Utrecht, NE), houses mounts for the emitters (2) and camera lenses (3). NIR VSMY1860 emitters (4) and ABS-printed lens replicas (5) are placed into these mounts. 3D-printed dummy lenses were used in the silicone casting process, allowing for easy removal without damaging the silicone layer. Those replicas are then replaced with actual lenses afterward.

After assembling all components, silicone (6) was prepared and deposited in the mold. Due to its well-known optical properties, particularly its optical transmittance, PDMS was chosen as the elastic element. A 30:1 ratio of PDMS base to curing agent was prepared and mixed using a Thinky Mixer (THINKY U.S.A., INC.), then degassed in a vacuum chamber for 7 m before being poured into the mold. This ratio was chosen to make the PDMS as soft as possible, enhancing the sensor's sensitivity and flexibility (Sales et al., 2021). This softer composition facilitates the detection of subtle interactions and deformations, improving its responsiveness

and the overall performance. Finally, the sensor was cured in an oven for 2 h at 75 c.

2.2 Experimental setup and data acquisition

Data acquisition in the setup shown in Figure 2 is carried out by applying an external force to the silicone layer using an Omega 3 Force Dimension robot. This leads to deformation and brightness changes detectable by the DVS cameras. The sensing area is rectangular, with DVS cameras at two corners. This arrangement covers most of the tactile sensing area, allowing each camera to capture variations in brightness caused by the deformation of the silicone layer illuminated by NIR LEDs placed on the other sides, as illustrated in Figure 1. The NIR LEDs (16 mW total power consumption) provide consistent illumination, thus ensuring high visibility of the silicone layer even in low-light conditions, enhancing the accuracy of touch-point detection. This configuration is expected to support efficient and highly dynamic range data acquisition, enabling responsive and adaptable sensing in robots and haptic settings.

The sensor was pressed at 250 locations arranged in a meander path with 4 mm spacing between each point at a constant depth of 2 mm (it should be mentioned that the applied force was not analyzed). To ensure data consistency and robustness, the grid pressing pattern was repeated 10 times, each involving the same 250 pressing points. After completing each set of 250 presses, the robot returns to the starting position to begin the next repetition.

To ensure temporal alignment between the two DVS recordings, a synchronization procedure is performed prior to the main experiment. Since the recordings were started manually and were not inherently synchronized, the center of the skin—clearly visible to both cameras—was pressed three times with 1-s intervals, followed by a 3-s pause before the first actual press. These three initial press actions serve as temporal landmarks, enabling a precise alignment of timestamps across both camera streams. The aligned event sequences ensure the accurate segmentation and a consistent comparison of press timings. These initial calibration actions are highlighted in Figure 4A with a rectangle.

The accumulated event rates for contacts on the circular path (48 presses on two circles with radii of 20 mm and 25 mm) are shown in Figure 3A. This dataset is used solely to illustrate how press-induced event activity appears on the sensor surface, as the circular pattern provides visually clear and evenly distributed examples of event generation. No localization analysis is performed on this dataset. All quantitative results presented in this work—including triangulation accuracy, pass-rate analysis, and detection latency distribution evaluation—are based exclusively on the meandering-path dataset shown in Figures 4, 5.

Much of the sensing area remains unused, as seen by the outer ring around the viewable area. The images in Figures 3B–D represent individual contacts on the same camera, each shown in a different color. These images highlight the varying activity levels across different pressure locations, illustrating how each press action generates a distinct pattern of event distribution on the sensor.

TABLE 1 Comparison of this work and similar tactile sensors.

Sensor	Hardware	Primitive	Algorithm	Metrics
Force-EvT (Guo et al., 2024)	DVS	Force estimation	Vision transform	RMSE 0.13 n (Newtons) $R^2 = 0.3$ 1.5 % error
Vision-based (Naeini et al., 2020)	DVS	Contact force	Time delay neural network	RMSE 0.16–0.17 n 7.2 % accuracy
Bio-inspired (Faris et al., 2024)	DVS	Slip & Pressure	Spike processing	5 % accuracy 2 ms RT
Evetac (Funk et al., 2024)	DVS	Vibration & Shear force	DL	Vibration sensing up to 48 h
E-BTS (Mukashev et al., 2025)	DVS	Force estimation	DL force prediction	RMSE ~ 1 n Latency < 5 ms
NeuroTac (Ward-Cherrier et al., 2020)	DVS	Texture/object recognition, slip cues	Spiking neural network (event-based)	Classification accuracy (task-dependent; e.g., 95% on Braille)
Multi-cam (Trueeb et al., 2020)	4 RGB cameras with silicone	Force distribution	DCNN	RMSE 0.057 n (Fz) 40 h
DIGIT (Lambeta et al., 2020)	RGB camera with elastomer	Shape & force	Autoencoder	640 × 480 px @ 60 h
Retrographic (Johnson and Adelson, 2009)	RGB with elastomer	Texture/shape	Photometric stereo	High-res 2.5D capture
MSOWS (Lo Preti et al., 2022b)	LED with photoreceiver	Localization	Sequential scan	5 mm resolution 2.6 % error
This work	Stereo DVS	Contact location	Stereo with DBSCAN	RMSE 4.66 mm 31 ms detection latency distribution

Each contact captured by the DVS cameras exhibits unique spatio-temporal characteristics, as illustrated in the sequence of images in Figures 3E–H on the right, which show the stream of events generated from a single press with events accumulated over a fixed time window. In these images, gray pixels represent areas without events, while black and white pixels indicate events with negative and positive polarities, respectively, corresponding to decreases and increases in light intensity detected by the sensor. This reflects how light from LEDs is blocked or revealed during skin deformation as the indenter presses and releases. As illustrated in Figure 2, the cameras are embedded at the edge of the skin, looking through the side of the silicone. The events are generated by the indenter pressing the top surface of the skin, causing deformations and interrupting the light coming from the LEDs visible from both cameras. This temporal progression of events highlights changes in the activity and could be leveraged for further processing, extracting additional information from the data, such as the classification of contact types, slip detection, or motion patterns. However, this study is focused on the localization of contacts based on aggregated event statistics.

2.3 Data processing and press actions segmentation

The DVS cameras generate continuous event data, with each press action on the optical skin generating a series of events. Initial data processing begins by calculating the event rate (number of events per second) across all press actions, offering a high-level overview of tactile activity over time. Figure 4A displays the event

rate from both Camera 1 (red) and Camera 2 (blue) over the first 100 s of the experiment. This corresponds to 27 individual pressures applied in a meandering path, as illustrated in the inset on the top left. The inset also shows the spatial layout of the sensor and the robot's pressing path, with the relative positions of the two cameras (red and blue) and the pressure trajectory over the optical skin. Each contact location is marked, and selected pressures are connected to their corresponding spikes in the event-rate plot using green rectangles for visual reference.

To remove irrelevant activity and focus only on tactile interactions, the upper and lower part of the visual field was cropped (on the v -axis, from $v = 200$ to $v = 360$), based on the 4 mm thickness of the silicone layer and the lens characteristics. The cropped region corresponds to the yellow box shown in Figure 3A.

Furthermore, to enable press-wise analysis, event data were segmented temporally into individual pressure events. Figure 4B zooms in on a single contact and shows the corresponding event-rate for Camera 1 (red) and Camera 2 (blue), highlighting synchronization and sensitivity of both views during a single tactile event.

2.4 Event detection and finding the centroids

After segmenting the individual pressures, we used the clustering-based approach Density-Based Spatial Clustering of Applications with Noise (DBSCAN) to find the horizontal centroid of each contact on the sensor of each camera.




FIGURE 1

Exploded view of the sensor design. Components: (1) ABS printed base; (2) emitter mounts; (3) lens mounts; (4) near-infrared VSMY1860 emitters; (5) ABS lens replicas used during molding; (6) PDMS tactile layer.

DBSCAN identifies the density of the clusters of contact activity without prior assumptions about the number of events or the shape of the clusters. As all presses are separated into individual containers based on their timing, DBSCAN can be applied to spatial data to find the region with maximum activity. Events that do not belong to any cluster are excluded as noise. DBSCAN was applied with an ϵ (eps) neighborhood of 10 pixels and a minimum of 10 events per cluster ($min_samples = 10$). These parameters were chosen empirically to balance noise rejection with robust cluster formation across different presses.


When one or more clusters are found, the largest cluster is computed and considered as the direction of the contact activity on the sensor's surface. This method filters out scattered noise events and focuses on the activity elicited by pressure contacts. From the remaining events, the centroid of the dominant cluster is computed as the mean of the (u, v) positions within the cluster. This centroid is used to represent the pressure location in the image space. Although both u - and v -coordinates are used for clustering, only the u -coordinates of the resulting centroids are retained for subsequent triangulation Section 2.6. This is because the two cameras are arranged along a horizontal stereo baseline, meaning that disparity, and consequently depth information, is primarily encoded in the horizontal (u) direction. The v -coordinate variations do not contribute to the triangulation accuracy in this configuration but remain useful for verifying the spatial consistency of the detected

contact regions. The process is repeated for all presses across both cameras.

Figures 5A, B present the top-down distribution of centroid u -positions identified by DBSCAN for contacts numbered from 80 to 150, visualized from the perspectives of Camera 1 and Camera 2, respectively. This visualization offers a mapping of detected pressure locations across the sensor surface. Additionally, Figures 5C, D show individual contact activity distributions with their clustered activity regions.

2.5 Data filtering and pressures exclusion criteria

Figure 5 illustrates scenarios in which press localization may fail due to weak or ambiguous event patterns in either view. Such weak pressure activities are highlighted by green circles in Figures 5A, B. To ensure a robust localization, a press was removed if DBSCAN failed to identify a prominent cluster of activity in one of the views. This happens when there is not enough density or structure in the spatial arrangement of events to create a significant cluster, which might indicate that the pressure was too weak or was obscured. A threshold number of cluster points was introduced as a metric to determine the existence of a valid pressure event. To perform the optimization in Section 2.6, some outliers were manually excluded



(this is consistent with a one-off calibration process for a new sensor). However, the outliers remaining after thresholding were included.

2.6 Geometric triangulation for pressure localization

Once the area of maximum activity for the u -coordinate of each camera is found, the angle between the center pixel of each camera's field of view and the identified active pixel (notated as θ_1 for Camera 1 and θ_2 for Camera 2) is computed. These angles, along with the known distances (that is, d_1 and d_2) from each camera to the center of the silicone layer, allow us to draw triangulation lines from both cameras toward the detected pressure location. Angles and distances are inputs for geometric triangulation, providing the basis to estimate the press location on the sensor's x - and y -coordinates.

The final pressure position is determined by computing the intersection of the triangulation lines from each camera, representing the estimated coordinates (x, y) of the press on the optical skin. DBSCAN allows for the estimation of the activity center, effectively mitigating noise and identifying center regions of pixel activity. Figure 6 illustrates this triangulation process, detailing the angles and distances of locating each press in the

sensing area. Although initial triangulation was performed using camera positions measured directly from the robot workspace, further refinement was introduced to improve localization accuracy by treating the following system parameters as optimization variables. Skew angles were introduced into the triangulation model to account for any possible misalignment between the cameras and the edges of the tactile surface. Recognizing that the cameras might not be perfectly parallel to the skin's edges due to setup imperfections, the skew angles and camera positions are treated as variables within the triangulation process. By incorporating skew angles (θ_{skew_1} for Camera 1 and θ_{skew_2} for Camera 2) and optimizing camera positions (x_1, y_1 for Camera 1 and x_2, y_2 for Camera 2) and lens distortion, the accuracy of estimated press locations was improved. The least squares optimization was then performed on these parameters to minimize the difference between the estimated pressure locations and the coordinates provided by the Omega 3 force dimension robot. This optimization was performed for the grid press trajectories to account for any differences in the data patterns. Optimized camera positions, skew angles, and distortion correction improved accuracy, as reflected in the error metrics presented in Section 3.1. It should be noted that, as described in the experimental setup, the grid press path consisting of 250 locations was repeated ten times. The triangulation parameters were optimized using data from one repetition of the path, and the optimized parameters were then applied to the remaining nine repetitions. This approach was adopted to minimize overfitting and to provide a more reliable estimate of the system's localization accuracy.

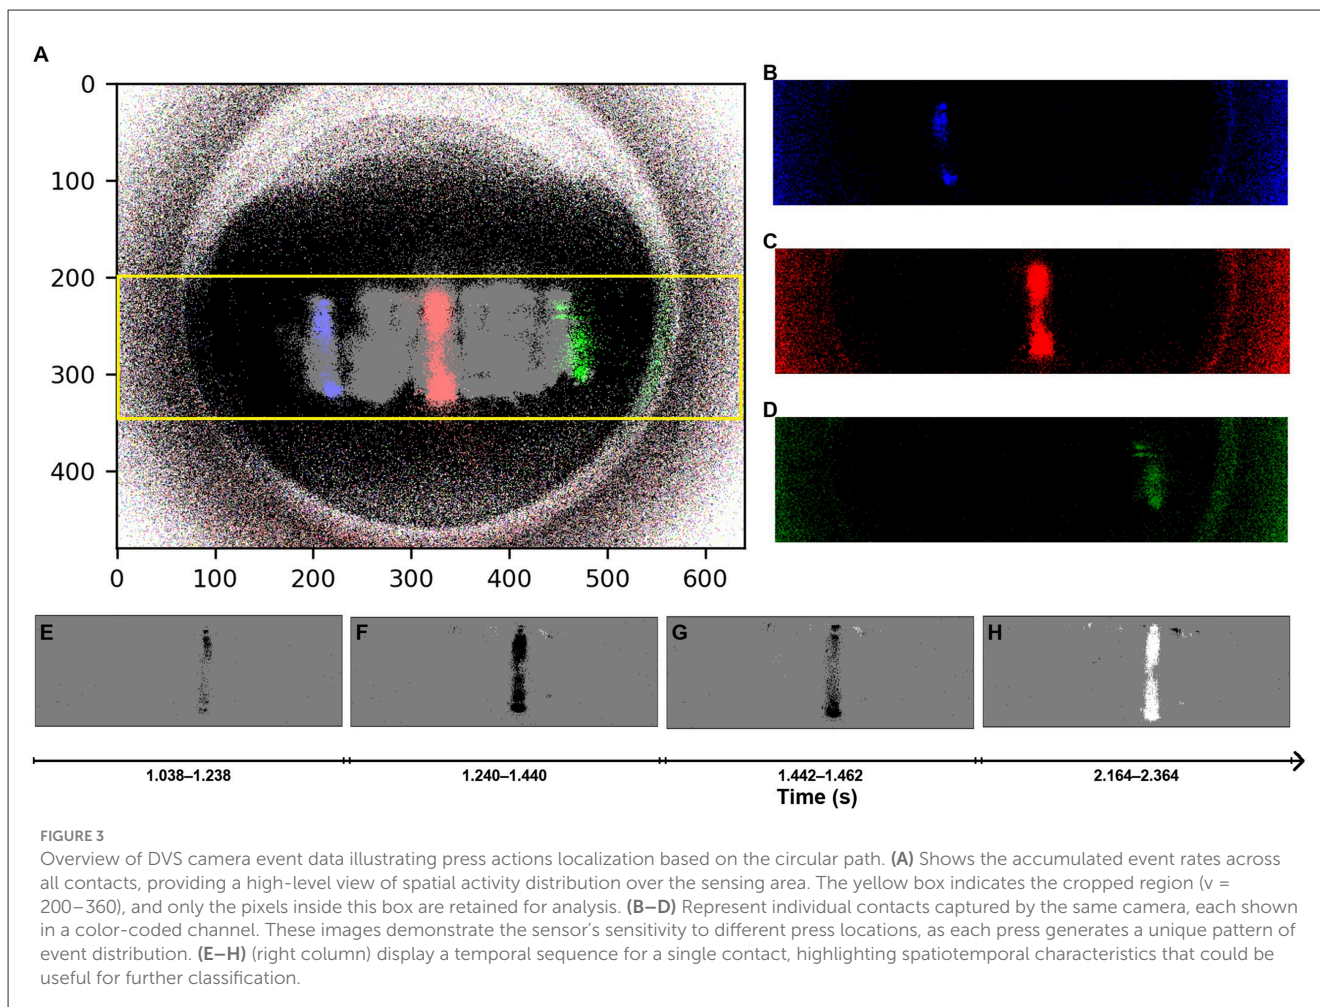
The differences in the optimized camera positions for the grid path are visible in Figure 7, where the camera positions are depicted differently for the path. Triangulation calculations are based on these introduced and optimized variables.

3 Results

This Section presents the performance of the proposed event-based optical skin in localizing contact events. We report the system's localization accuracy based on ground truth comparisons, analyze its spatial coverage and detection robustness, and evaluate its response to data reduction through ablation testing. In addition, the potential latency distribution of tactile event detection is quantified to assess the sensor's suitability for real-time applications.

3.1 Error calculation and comparison

To evaluate the accuracy of the triangulation method, the estimated pressure locations are compared to the ground-truth coordinates provided by the Omega 3 force dimension robot. The Root-Mean-Squared Errors (RMSE) was selected as the primary performance indicator among the calculated error metrics. The RMSE for the grid path is 4.66 mm, which corresponds to 3.75 % of the sensor's diagonal. The mean standard deviation between trials for single pressures on the grid path is 2.85 mm. This indicates that the triangulation approach can estimate press positions, providing sufficient accuracy for practical tactile sensing



applications. Figure 7 illustrates the estimated pressure locations derived from DBSCAN for the grid path. To better understand how the localization error is distributed, we also calculated the RMSE separately along the x - and y -coordinates. The RMSE in the x -direction is 4.17 mm, and the RMSE in the y -direction is 3.28 mm. These values show how the triangulation geometry contributes differently along each axis and complement the overall Euclidean RMSE of 4.66 mm.

The total area of the fabricated skin surface is 9,328 mm. Out of this area, 8,763 mm (93.9 %) is within the field of view of both cameras. The robot probed 4,620 mm (52.7 %) of the visible skin area. After applying a manual threshold to reject poor quality results, the proportion of probed locations, which are visible to both cameras, where pressures can be reliably detected and localized, is 95 %.

3.2 Data ablation and latency distribution analysis

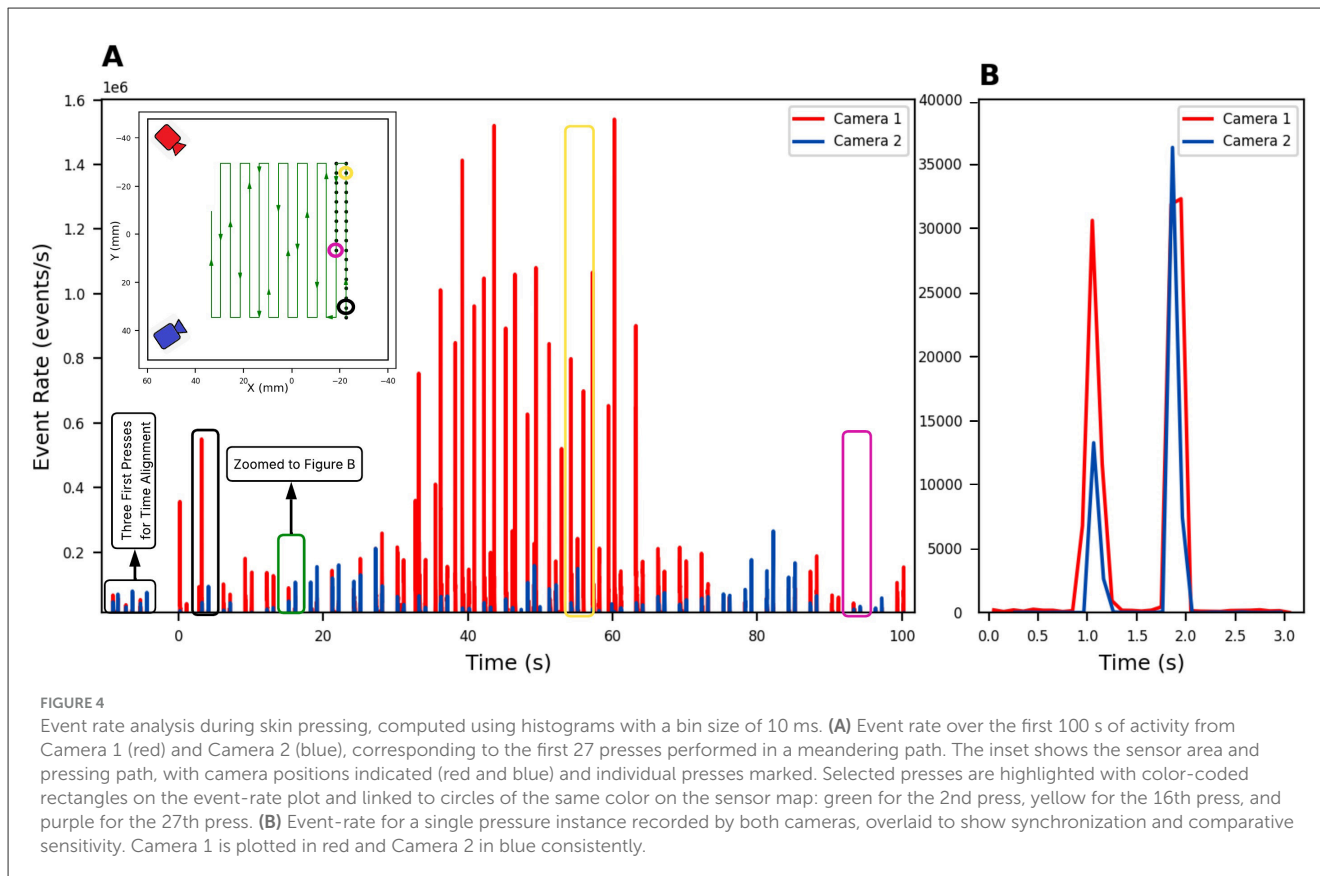
One of the main motivations for this work is achieving a low-data sensor maintaining the accuracy of the localization. High-resolution DVS cameras are utilized in the present realization.

Subsequent system versions could use dedicated event sensors with substantially fewer pixels, consequently reducing data transmission and processing demands. Furthermore, optimization of sensor tuning would also enhance responsiveness, with better control over the event rate.

To investigate how much the system depends on the amount of event data, we performed a data ablation (event reduction) analysis. Event data was progressively reduced using predefined downsampling factors (2^0 to 2^{10}), and the localization accuracy was evaluated using the RMSE and pass rate metrics.

The downsampling was done by applying a simple stochastic thinning of the event stream. For a given reduction factor k , each event was kept with probability $1/k$ and discarded otherwise. This keeps the original timing of the events but reduces their total number by roughly a factor of k . For each reduction factor, we repeated the thinning process using different random seeds to check the consistency of the results.

For each press, we computed the Euclidean distance between the estimated and actual press coordinates. For a given reduction factor, the distribution of these errors was obtained over all presses. A press was counted as a “pass” if its localization error was smaller than the 95th percentile of that error distribution. The pass rate is the proportion of presses that meet this criterion. This gives an intuitive measure of how often the system



still produces a usable localization after reducing the number of events.

Figure 8 shows the pass rate as a function of the reduction factor. The pass rate stays above 85% for reduction factors up to 2^{10} (1,024), but drops beyond this point, indicating that too much event removal starts to eliminate essential information for localization.

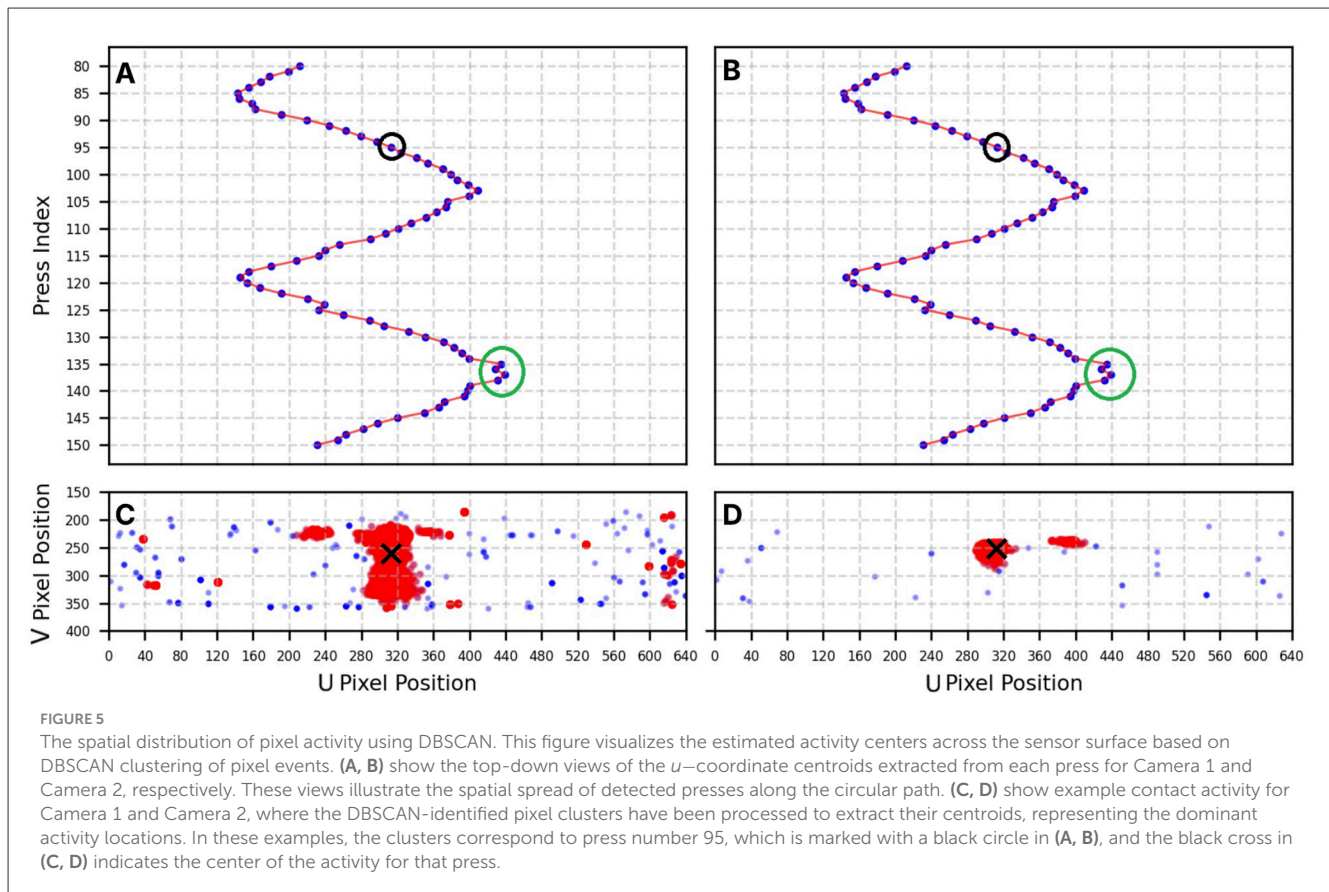
DVS are inherently fast due to their asynchronous event-driven nature, allowing near-instantaneous response to changes in the visual field (Delbruck et al., 2008; Posch et al., 2014). However, we evaluated the potential latency of this system in detecting a press by analyzing event-rate distributions over time, to help understand whether the low-latency of the DVS can be exploited.

Although the robot triggers were regular at known intervals, no hardware-recorded ground-truth onset is available. Therefore, we quantified detection latency distribution by finding the times when the aggregate event rates rise above normal noise and reporting the interval in which most of these crossings occur. For each trial, events were binned at high temporal resolution (0.2 ms bins) and smoothed with a short Gaussian kernel ($\sigma = 0.5$ ms). The baseline firing rate was estimated from a short background snippet immediately preceding the stimulus. Onset was then defined as the first threshold crossing of a one-sided Cumulative Sum (CUSUM) detector tuned to a four-fold rate increase over baseline, requiring at least three consecutive bins above threshold. A threshold parameter h was optimized by Receiver Operating Characteristic (ROC) analysis to ensure $\geq 95\%$ of trials were detected within a ± 100 ms window of the population median onset, while minimizing the false-alarm rate background

snippets.¹ Figure 9 shows 95th, 50th, and 5th percentile event rates for press versus no-press, with time aligned to the median onset time during presses ($h=1$). The detection latency distribution, i.e., the time between 5th and 95th percentile onset times for true positives, is 31 ms and the false-alarm rate is 0.13 events/s (where the 31 ms is used as a cool-down between consecutive false alarms).

This window is likely an overestimate of the latency from event production to the possibility of contact detection because the method only looks at event rates combined from both cameras and does not consider any other spatiotemporal cues, which could improve detection. At the same time, it underestimates the end-to-end system latency, since any delay from the robotic trigger to the production of the first press-related events is not included. Furthermore, this method is an analysis of the possibility of contact detection based on event statistics, rather than a proposed online algorithm. All of our processing in this work has been offline and trial-based and we have not presented an online method for press detection, therefore the detection latency distribution is also an underestimate of full-system latency for this reason. With 1,024× data reduction, 88% True Positive Rate (TPR) was achieved at $h = 3$, the latency distribution was 113 ms with 1.86 false-alarms per second.

¹ https://github.com/event-driven-robotics/optoskin/blob/main/latency/latency_from_rates_main.py



4 Discussion

In contrast to systems such as Chorley et al. (2009) and Kamiyama et al. (2004), and in common with MSOWS (Lo Preti et al., 2022b), this system is not dependent on marker tracking. Instead, it observes the effect of surface deformation on the passage of light through the transparent waveguide.

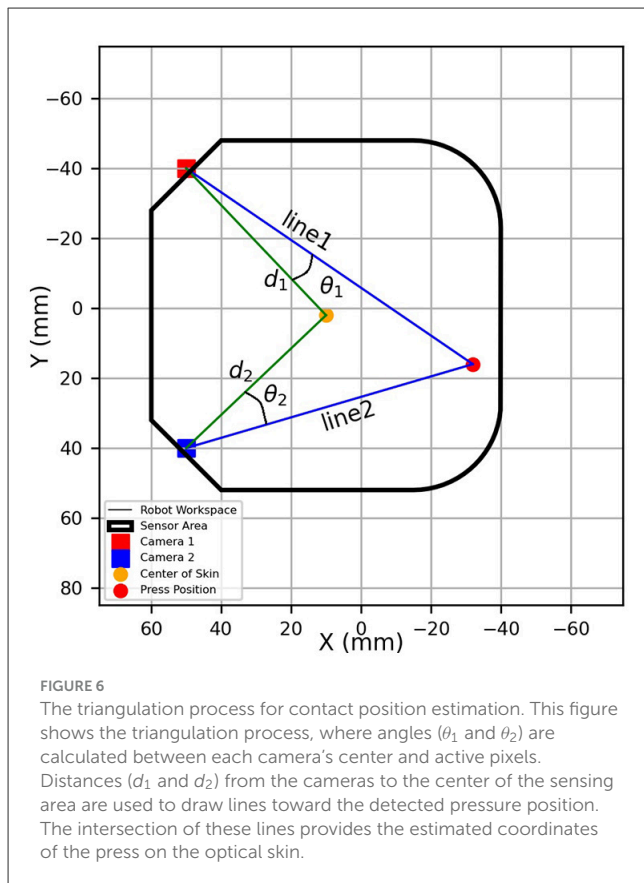
A comparison with MSOWS is given in Table 2. MSOWS uses a number of infrared emitters (PEs) and detectors (PRs) embedded all around the skin's periphery (24 of each). The PEs are strobed in sequence, the combinations of PEs and PRs define rays across the skin whose precision is limited by emission and reception cones, and the contact is localized by mapping the resulting reductions in received light at each PR to a virtual grid across the skin area (that is, a 16×16 grid with 5 mm cell pitch). In contrast, our DVS-based skin uses two 640×480 event-driven cameras with lenses, in combination with a number of PEs (32) around the skin periphery. The PEs are constantly illuminated so that events are not produced in static conditions. Therefore, the power cost for illumination is elevated in our system, however, synchronization of strobing is not necessary. The strobing rate places a limit on the temporal precision of MSOWS, which can only be improved by increasing the strobing rate and thus the data rate. The DVS-based system does not have such a restriction, since event cameras only produce events in the case of change. The latency of DVS is known to vary with illumination and can be as low as a few microseconds

with good illumination. The experimental scenario and the system presented here do not allow for the potential latency advantage to be demonstrated or evaluated, since a press can be reliably distinguished from background noise after 31 ms.

Given static illumination, the lenses are necessary to disambiguate the directions from which presses are detected and offer high precision for localization. In the method presented here, DBSCAN is used to cluster the events corresponding to a contact. The center of the detected cluster is extracted along the horizontal (u) axis of each camera and used to determine the angle of incidence based on the camera's field of view, which may, in principle, give an arbitrarily better accuracy.

The characterization of the localization accuracy is given in terms of the distance between the estimated location and the actual press location, with an RMSE of 4.66 mm. To contextualize this, an effective number of taxels can be estimated by calculating how many circles of radius 4.66 mm fit into the sensor area. Across the probed region of the sensor, this corresponds to 46 effective taxels (if the same RMSE were to hold across the full sensor area, this would be equivalent to 105 taxels, but we do not claim this, since the RMSE has only been measured within the probed region. It is likely that other effects, such as distortion or the proximity of the LEDs, would be more prominent toward the edges of the sensor).

For a better comparison with MSOWS, the Contact Map Reconstruction Error (CMRE) statistic for global error in a single press scenario (Lo Preti et al., 2022b) has been computed. This system achieves a global error of 3.75 % vs. MSOWS's 1.8 %.



The employed DVS cameras are not optimized for this application and may be highly overspecified. It may be that small pods of DVS pixels in 2 or more locations could deliver more efficient performance. To evaluate the possibility of using fewer event data, events have been randomly eliminated within trials, and the results shown in the Table 2 are for a 1,024 \times reduction. In this case, the CMRE rises to 6.4 %, RMSE rises to 9.33 mm, the effective number of taxels reduces to 34 (15 probed), and pressure detection latency distribution is approximately 31 ms and remains at about 113 ms under a 1,024 \times data reduction. By applying a 1024 \times data reduction, the DVS version can cut its peak press-mode bit-rate from 75 kB/s down to 1 kB/s, and its idle bit-rate from 24 kB/s to 0.02 kB/s, compared to MSOWS's constant 145 kB/s. MSOWS used 12-bit A-D conversion, whereas we have considered 21 bits per DVS event, given the size of the VGA address space.

In addition to comparing our system with MSOWS, we place it alongside other recent tactile localization technologies. These include resistive and tomographic skins, capacitive arrays, optical tactile sensors, magnetic skins, inductive and microwave sensors, and acoustic approaches. These systems vary widely in materials, wiring complexity, sensing principles, and achievable accuracy, but all aim to localize contact across a surface. The main differences between our DVS-based skin and these technologies concern sensing modality and data rate: our approach uses two event cameras instead of dense arrays, avoids global inverse problem solving, and naturally supports low idle bit-rate operation. A summary comparison of these technologies is provided in Table 3.

Only single-touch localization has been investigated in this work. Low-pressure detection and multi-touch interactions were not studied. Early tests (not shown here) agree with previous findings in MSOWS that contact can still be detected through internal light reflection when the surface is bent. With the reduced viewing locations provided by focused optics, the present system could also be extended to larger or more complex surfaces. Beyond localization, we did not attempt any further tasks such as identifying the type of contact or classifying the source of the pressure. However, the event data contain richer information than just position. Features such as the shape of the event cluster, the number of events produced, and the timing of the event activity could support future work on force estimation, slip detection, object recognition, or classification of different touch types. These tasks would require additional processing and calibration and remain outside the scope of this proof-of-concept study, but they represent clear next steps.

4.1 Limitations and future directions for multi-touch sensing

A key limitation of the current work is its exclusive focus on single-point contact localization. The methodology presented, which relies on identifying a single dominant event cluster in each camera, is fundamentally unequipped to resolve multiple simultaneous touch events. This is a critical consideration for future applications in complex, interactive scenarios.

Another limitation is that the current system does not estimate the magnitude of the applied pressure. This work was designed as a proof of concept focused on press localization. The robot pressed the sensor surface with a cubic tip to a fixed indentation depth of 2 mm; position control was used and the applied force was not measured during the experiments. However, the event data suggest that pressure estimation could be added in future work. As shown in Figures 6C, D, a single press produces a characteristic cluster of events on each camera. By analyzing features such as the size of the activated region, the density of events, or the temporal evolution of the cluster during indentation, it may be possible to infer the applied force. A dedicated calibration procedure would be required to map these features to physical pressure values.

Another limitation of the present work is that the evaluation was carried out only within the central region of the sensor surface. In early tests, presses near the edges of the silicone produced noticeably larger localization errors, as the image distortion increases toward the periphery. More importantly, the silicone layer had to be made very soft (30:1 PDMS, see Section 2.1) to maximize sensitivity, and this made the material mechanically fragile at the edges. Since the NIR LEDs are embedded directly in the silicone along the boundaries, pressing too close to an edge caused the silicone to lift and risk damage to the LEDs or tearing of the material. For these reasons, the robot was constrained to press only within the safer central region of the sensor during the experiments, corresponding to 52.7% of the usable sensing area. A more robust fabrication process, would be needed to allow full-area testing in future work.

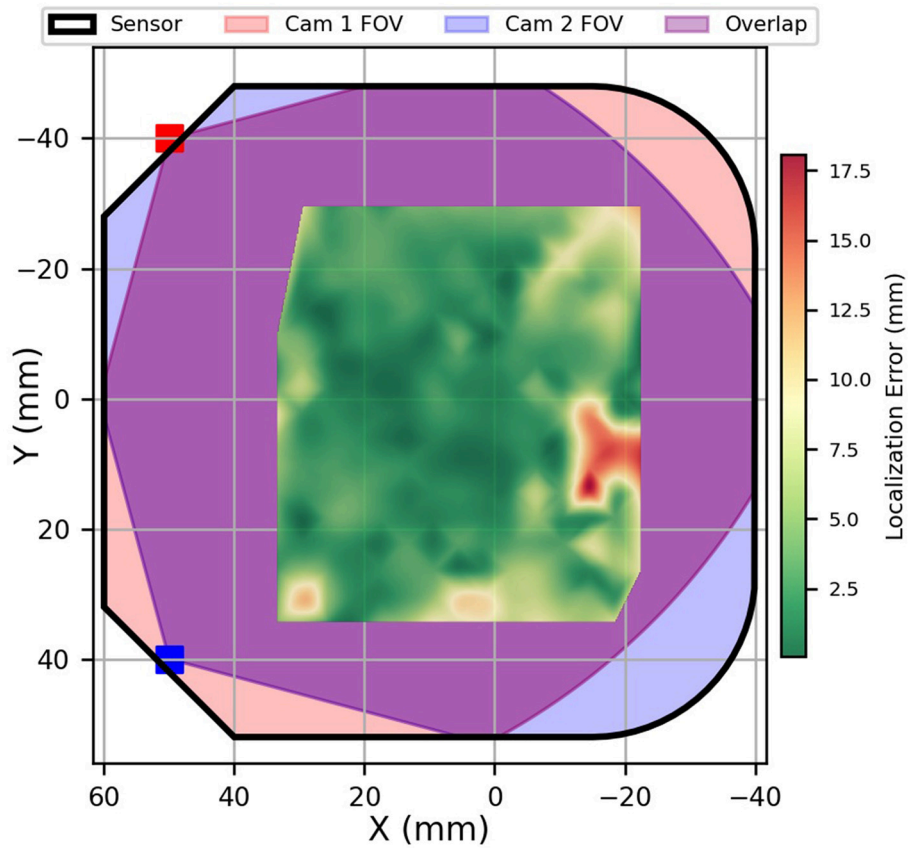


FIGURE 7 Localization error heatmap for the DBSCAN-based triangulation method on the grid path. Red regions indicate higher error magnitudes, while green regions correspond to more accurate estimations. The field of view (FOV) of each camera, their overlap region, and the sensor boundary are also shown.




FIGURE 8 Pass rate trend with event reduction. The system maintains a high detection success rate, even with a reduction factor of 1,024x.

The failure arises from the classic stereo correspondence problem. In a multi-touch scenario, each DVS camera would detect multiple, distinct clusters of events. The current algorithm would either erroneously merge nearby clusters into a single, inaccurate centroid or, if modified to detect multiple centroids,

would lack a mechanism to correctly pair the corresponding centroids from the left and right cameras. For instance, with two presses generating centroids L_1 and L_2 in the left camera's view and R_1 and R_2 in the right's, the system could not disambiguate between the correct pairings (L_1-R_1 , L_2-R_2)

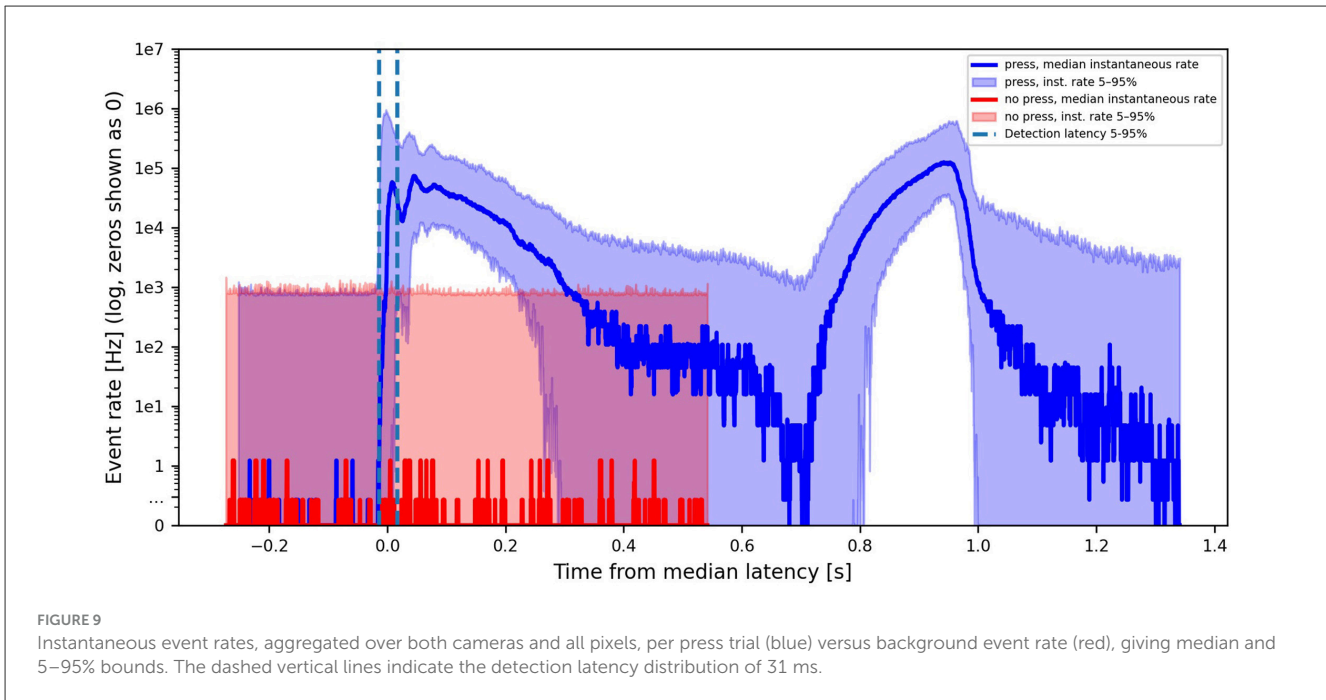


TABLE 2 Comparison of this work and MSOWS.

Metric	Lo Preti et al. (2022) (MSOWS)	This work (DVS-based)	
		No data reduction	1,024× reduction
Sensor type	Optical waveguide (PE-PR)	DVS (stereo)	
Data reduction	N/A	No data reduction	1024×
Number of view points	24	2	
Processing unit	FPGA DAQ board	On-board event processing	
Samples per press	168 samples/PR	Event stream (sparse)	
ADC bit depth	12-b	21-bit	
Press duration	0.144 s	0.55 s	0.55 s
Bit rate (press)	~145 kBs	~975 kBs	~1 kBs
Bit rate (idle)	~145 kBs	~24 kBs	~0.02 kBs
System latency	—	0.031 s Detection latency distribution	0.113 s Detection latency distribution
Global error	1.8 %	3.75 %	~6.4 %
RMSE (localization)	—	4.66 mm	9.33 mm
Effective num taxels over sensor area	—	105	34
Effective num taxels over probed area	—	46	15

and the incorrect “ghost” pairings (L_1-R_2 , L_2-R_1), leading to erroneous localization.

However, the event-based nature of the DVS offers a powerful solution. It is highly improbable that two distinct mechanical presses would be initiated in the exact same microsecond. Future work can implement a temporal correlation filter. By segmenting event clusters based on their onset timestamps, the system can pair clusters from the two cameras that appear and evolve within a narrow, shared time window (e.g., ≤ 1 ms), thus robustly solving the correspondence problem for non-synchronous contacts.

An alternative approach is to move beyond simple centroiding and utilize richer spatial features of the event clusters. The shape, size, total event count, and velocity profile of corresponding clusters should be similar across both camera views. A future algorithm could combine epipolar geometry constraints with a feature-matching technique (e.g., comparing cluster moments or event density histograms) to calculate a matching score and identify the most probable pairings.

For guaranteed disambiguation in cluttered or highly dynamic scenarios, the system could be extended with a third, non-collinear

TABLE 3 Comparison of representative tactile localization technologies.

Sensor	Localization capabilities	Materials	Accuracy/resolution	Techniques (hardware/algorithm)
EIT sheet (Chen et al., 2025b)	Multi-touch with bending	Conductive hydrogel; elastomer layer	5.4 mm error	EIT reconstruction with ML for touch-bend decoupling
Hydrogel EIT skin (Hardman et al., 2023)	Multi-touch	Self-healing ionic hydrogel	12.1 mm	Data-driven EIT reconstruction
Stretchable capacitive array (Sarwar et al., 2017)	Multi-touch under stretch/bending	Transparent elastomer; stretchable electrodes	1 cm resolution	Capacitive grid with deformable substrate
Magnetic skin (Hu et al., 2024)	Multi-touch, multi-scale; large-area sensing	Magnetized elastomer film; hall sensors	1.2 mm error over 48,400 mm area	Signal processing with CNN
Graphene Hall magnetic array (Li et al., 2025)	Single-touch with virtual 6×6 resolution	Graphene Hall elements; magnetized layer	1.3 mm average error	Vertical periodic magnetization to reduce field interference
Flexible EIT tactile skin (Chen et al., 2025a)	Multi-touch with bending	Magnetic hydrogel layer and Ecoflex	4.8 ± 2.8 mm touch error; bend RMSE ~0.9°	CNN state classifier
SonicBoom (Lee et al., 2025)	Single-touch	MEMS microphones	0.43–2.22 cm localization error	Vibration-based triangulation with learned model

camera. This would provide an additional geometric constraint, allowing for robust triangulation even with ambiguous pairings between any two of the views. The intersection point validated by all three viewpoints would confirm a true contact, while 'ghost' points would be rejected for failing this multi-view consistency check.

Beyond the current triangulation-based method, an alternative approach to resolving the stereo correspondence problem in DVS-based tactile sensing is the use of Spiking Neural Networks (SNNs). Osswald et al. (2017) proposed an SNN model that enables 3D perception by leveraging spike-based representations and temporal dynamics.

5 Conclusions

This work presents an optical skin system designed for advanced tactile sensing using DVS cameras with a compliant silicone layer. Based on the MSOWS design (Lo Preti et al., 2022b), this work proposes a marker-free event-based tactile sensor by replacing embedded photoreceivers with DVS cameras. Our approach allows for localization of contacts through stereo vision and triangulation, with a simple, flexible, and responsive design suitable for dynamic tactile applications. The system achieves a pressure localization accuracy of 4.66 mm RMSE. Data reduction analysis reveals that the event-based approach could sustain a pass rate above 85 % even when reducing data transmission by a factor of 1,024, confirming the feasibility of a low-bandwidth implementation. The detection latency distribution is 31 ms, demonstrating the potential system's suitability for real-time tactile interactions.

By utilizing DVS technology, the system benefits from high temporal resolution and reduced data bandwidth, making it well-suited for scenarios requiring real-time processing and high-speed interaction. Furthermore, the system's marker-free design simplifies integration, paving the way for scalable and adaptable soft robotic applications.

6 Future work

While the current design is a proof-of-concept that successfully localizes contacts, several areas for improvement remain.

Future work will focus on a new design that grants the optical skin higher sensitivity. The next envisioned steps include extensive testing on various tactile tasks, such as precise localization, slip detection, and force mapping. These would be useful in fine-tuning the system to achieve better accuracy and reliability, and become more easily deployable in soft robots and interactive objects. Unused data features such as V-address, polarity, and precise timing may be used to increase the system's discrimination ability. Optics may be better matched to the employed cameras. Smaller cameras with fewer pixels may provide acceptable performance with even lower data rates. Moving toward a design with small clusters of pixels and dedicated optics introduces key trade-offs. While it reduces the overall data processing load and lowers bandwidth requirements, it increases wiring complexity. Additional synchronization strategies may be required between clusters to maintain localization accuracy. Furthermore, decreasing the number of pixels limits the resolution of event detection, potentially impacting fine-grained spatial accuracy. However, a well-optimized optical system could compensate for these constraints, enabling a balance between data efficiency and sensing performance. SNN-based methods might be investigated as an alternative to triangulation, potentially improving accuracy and resolving remaining systematic skew and adaptability to more complex tactile interactions.

Data availability statement

The dataset used in this study is available at Zenodo: <https://zenodo.org/records/17274644>. The code for processing and analysis is available on GitHub: <https://github.com/event-driven-robotics/optoskin>.

Author contributions

MK: Methodology, Visualization, Data curation, Software, Conceptualization, Writing – review & editing, Formal analysis, Validation, Writing – original draft. SB: Validation, Methodology, Visualization, Supervision, Writing – original draft, Writing – review & editing, Software. PT: Writing – original draft, Conceptualization, Writing – review & editing. SM-C: Writing – original draft, Data curation, Writing – review & editing. ML: Writing – review & editing, Writing – original draft. FM: Writing – original draft, Writing – review & editing. LB: Writing – original draft, Writing – review & editing. CB: Writing – review & editing, Writing – original draft, Investigation, Supervision, Funding acquisition.

Funding

The author(s) declared that financial support was received for this work and/or its publication. CB, SB, MK, PT, and LB acknowledge the support of the IIT Flagship program Brain & Machines and Technologies for Sustainability. CB acknowledges the financial support of the National Biodiversity Future Center funded under the National Recovery and Resilience Plan (NRRP), Mission 4 Component 2 Investment 1.4 - Call for tender No. 3138 of 16 December 2021, rectified by Decree n.3175 of 18 December 2021 of Italian Ministry of University and Research funded by the European Union – NextGenerationEU. SB acknowledges the financial support from PNRR MUR Project PE000013 “Future Artificial Intelligence Research (hereafter FAIR)” funded by the European Union – NextGenerationEU. SB acknowledges the

support of the European Union Horizon Europe programme, grant agreement ID: 101120727, PRIMI: Performance in Robots Interaction via Mental Imagery.

Conflict of interest

The author(s) declared that this work was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declared that generative AI was not used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

- Bartolozzi, C., Natale, L., and Nori, F. (2016). Robots with a sense of touch. *Nat. Mater.* 15, 921–925. doi: 10.1038/nmat4731
- Brandli, C., Berner, R., Yang, M., Liu, S.-C., and Delbruck, T. (2014). A 240 × 180 130 db 3 μ s latency global shutter spatiotemporal vision sensor. *IEEE J. Solid-State Circ.* 49, 2333–2341. doi: 10.1109/JSSC.2014.2342715
- Chen, H., Himmel, B., Li, B., Wang, X., and Hoffmann, M. (2025a). Multi-touch and bending perception using electrical impedance tomography for robotics. *arXiv preprint arXiv:2503.13048*.
- Chen, Z., Zhang, Y., Wang, J., and Gao, H. (2025b). Multi-touch and bending perception using electrical impedance tomography for robotics. *arXiv preprint arXiv:2503.13048*.
- Chorley, C., Melhuish, C., Pipe, T., and Rossiter, J. (2009). “Development of a tactile sensor based on biologically inspired edge encoding,” in *2009 International Conference on Advanced Robotics* (IEEE), 1–6.
- Delbruck, T. (2008). “Frame-free dynamic digital vision,” in *Proceedings of Intl. Symposium on Secure-Life Electronics, Advanced Electronics for Quality Life and Society*, 21–26.
- Faris, O., et al. (2023). Proprioception and exteroception of a soft robotic finger using neuromorphic vision-based sensing. *Soft Robot.* 10, 467–481. doi: 10.1089/soro.2022.0030
- Faris, O., Awad, M. I., Awad, M. A., Zweiri, Y. H., and Khalaf, K. (2024). A novel bioinspired neuromorphic vision-based tactile sensor for fast tactile perception. *arXiv preprint, abs/2403.10120*.
- Funk, N., Helmut, E., Chalvatzaki, G., Calandra, R., and Peters, J. (2024). Evetac: an event-based optical tactile sensor for robotic manipulation. *IEEE Trans. Robot.* 40, 3812–3832. doi: 10.1109/TRO.2024.3428430
- Gallego, G., Delbruck, T., Orchard, G., Bartolozzi, C., Taba, B., Censi, A., et al. (2022). Event-based vision: a survey. *IEEE Trans. Pattern Anal. Mach. Intell.* 44, 154–180. doi: 10.1109/TPAMI.2020.3008413
- Guo, Q., Yu, Z., Fu, J., Lu, Y., Zweiri, Y., and Gan, D. (2024). Force-evt: a closer look at robotic gripper force measurement with event-based vision transformer. *arXiv preprint, arXiv:2404.01170*.
- Guo, X. H., Hong, W. Q., Liu, L., Wang, D. D., Xiang, L., Mai, Z. H., et al. (2022). Highly sensitive and wide-range flexible bionic tactile sensors inspired by the octopus sucker structure. *ACS Appl. Nano Mater.* 5, 11028–11036. doi: 10.1021/acsnm.2c02242
- Hardman, D., Thuruthel, T. G., and Iida, F. (2023). Tactile perception in hydrogel-based robotic skins using data-driven electrical impedance tomography. *Mater. Today Electr.* 4:100032. doi: 10.1016/j.mtelec.2023.100032
- Hu, H., Zhang, C., Lai, X., Dai, H., Pan, C., Sun, H., et al. (2024). Large-area magnetic skin for multi-point and multi-scale tactile sensing with super-resolution. *NPJ Flexible Electr.* 8:42. doi: 10.1038/s41528-024-00325-z
- Johnson, M. K., and Adelson, E. H. (2009). “Retinographic sensing for the measurement of surface texture and shape,” in *Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition*. doi: 10.1109/CVPRW.2009.5206534
- Kamiyama, K., Kajimoto, H., Kawakami, N., and Tachi, S. (2004). “Evaluation of a vision-based tactile sensor,” in *Proceedings of the 2004 IEEE International Conference on Robotics and Automation* (IEEE). doi: 10.1109/ROBOT.2004.1308043
- Kang, K., Kim, K., Baek, J., Lee, D. J., and Yu, K. (2024). Biomimic and bioinspired soft neuromorphic tactile sensory system. *Appl. Phys. Rev.* 11, 1–30. doi: 10.1063/5.0204104
- Lambeta, M., Chou, P. W., Tian, S., Yang, B., Maloon, B., Most, V. R., et al. (2020). Digit: a novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation. *IEEE Robot. Autom. Lett.* 5, 3838–3845. doi: 10.1109/LRA.2020.2977257

- Lambeta, M., Wu, T., Sengul, A., Most, V. R., Black, N., Sawyer, K., et al. (2024). Digitizing touch with an artificial multimodal fingertip. *arXiv preprint arXiv:2411.02479*.
- Lee, M., Yoo, U., Oh, J., Ichnowski, J., Kantor, G., and Kroemer, O. (2024). Sonicboom: contact localization using array of microphones. *arXiv preprint, arXiv:2412.09878*.
- Lee, M., Yoo, U., Oh, J., Ichnowski, J., Kantor, G., and Kroemer, O. (2025). Sonicboom: contact localization using array of microphones. *IEEE Robot. Autom. Lett.* 10, 7603–7610. doi: 10.1109/LRA.2025.3576067
- Lepora, N. F. (2021). Soft biomimetic optical tactile sensing with the TACTIP: a review. *IEEE Sens. J.* 21, 21131–21143. doi: 10.1109/JSEN.2021.3100645
- Li, X., Gao, W., Wang, B., Jiao, W., Su, H., Zhou, W., et al. (2025). Graphene-based flexible magnetic tactile sensor with vertically periodic magnetization for enhanced spatial resolution. *Soft Sci.* 5:58. doi: 10.20517/ss.2025.56
- Li, Y., Hu, J., Cao, D., Wang, S., Dasgupta, P., and Liu, H. (2022). Optical-waveguide based tactile sensing for surgical instruments of minimally invasive surgery. *Front. Robot. AI* 8:773166. doi: 10.3389/frobt.2021.773166
- Lichtsteiner, P., Posch, C., and Delbruck, T. (2008). A 128×128 120 db 15 μ s latency asynchronous temporal contrast vision sensor. *IEEE J. Solid-State Circ.* 43, 566–576. doi: 10.1109/JSSC.2007.914337
- Liu, T., Brayshaw, G., Li, A., Xu, X., and Ward-Cherrier, B. (2025). Neuromorphic touch for robotics—a review. *Neuromor. Comput. Eng.* 5, 1–25. doi: 10.1088/2634-4386/adf091
- Lo Preti, M., Totaro, M., Falotico, E., and Beccai, L. (2022a). “Optical-based technologies for artificial soft tactile sensing,” in *Electronic Skin* (River Publishers), 73–99. doi: 10.1201/9781003338062-4
- Lo Preti, M., Totaro, M., Falotico, E., Crepaldi, M., and Beccai, L. (2022b). Online pressure map reconstruction in a multitouch soft optical waveguide skin. *IEEE/ASME Trans. Mechatr.* 27, 4530–4540. doi: 10.1109/TMECH.2022.3158979
- Mukashev, D., Seitzhan, S., Chumakov, J., Khajikhanov, S., Yergibay, M., Zhaniyar, N., et al. (2025). E-bts: event-based tactile sensor for haptic teleoperation in augmented reality. *IEEE Trans. Robot.* 41, 450–463. doi: 10.1109/TRO.2024.3502215
- Naeini, F. B., AlAli, A. M., Al-Husari, R., Rigi, A., Al-Sharman, M. K., Makris, D., et al. (2020). A novel dynamic-vision-based approach for tactile sensing applications. *IEEE Trans. Instrum. Meas.* 69, 1881–1893. doi: 10.1109/TIM.2019.2919354
- Osswald, M., Ieng, S. H., Benosman, R., Indiveri, G. (2017). A spiking neural network model of 3D perception for event-based neuromorphic stereo vision systems. *Sci. Rep.* 7:40703. doi: 10.1038/srep40703
- Posch, C., Serrano-Gotarredona, T., Linares-Barranco, B., and Delbruck, T. (2014). Retinomorphic event-based vision sensors: bioinspired cameras with spiking output. *Proc. IEEE* 102, 1470–1484. doi: 10.1109/JPROC.2014.2346153
- Preti, M. L., Trunin, P., Cafiso, D., and Beccai, L. (2025). “The hexsows: a modular soft and robust skin concept for multimodal tactile sensing,” in *2025 IEEE 8th International Conference on Soft Robotics (RoboSoft)* (Lausanne, Switzerland), 1–8. doi: 10.1109/RoboSoft63089.2025.11020900
- Rigi, A., Naeini, F. B., Makris, D., and Zweiri, Y. (2018). A novel event-based incipient slip detection using dynamic active-pixel vision sensor (Davis). *Sensors* 18:333. doi: 10.3390/s18020333
- Sales, F., Souza, A., Ariati, R., Noronha, V., Giovanetti, E., Lima, R., et al. (2021). Composite material of PDMS with interchangeable transmittance: study of optical, mechanical properties and wettability. *J. Comp. Sci.* 5:110. doi: 10.3390/jcs5040110
- Sankar, S., Cheng, W.-Y., Zhang, J., Slepian, A., Iskarous, M. M., Greene, R. J., et al. (2025). A natural biomimetic prosthetic hand with neuromorphic tactile sensing for precise and compliant grasping. *Sci. Adv.* 11:eadr9300. doi: 10.1126/sciadv.adr9300
- Sarwar, M. S., Dobashi, Y., Preston, D. J., Shea, H., Wood, R. J., and Walsh, C. J. (2017). Bending and stretching of a highly stretchable transparent electronic skin. *Sci. Adv.* 3:1602200. doi: 10.1126/sciadv.1602200
- Sferrazza, C., and D’Andrea, R. (2019). Design, motivation and evaluation of a full-resolution optical tactile sensor. *Sensors* 19:928. doi: 10.3390/s19040928
- Shimonomura, K. (2019). Tactile image sensors employing camera: a review. *Sensors* 19:3933. doi: 10.3390/s19183933
- Trueeb, C., Sferrazza, C., and D’Andrea, R. (2020). “Towards vision-based robotic skins: a data-driven, multi-camera tactile sensor,” in *2020 3rd IEEE International Conference on Soft Robotics (RoboSoft)*, 333–338. doi: 10.1109/RoboSoft48309.2020.9116060
- Trunin, P., Cafiso, D., and Beccai, L. (2025). Design and 3d printing of soft optical waveguides towards monolithic perceptive systems. *Addit. Manuf.* 100:104687. doi: 10.1016/j.addma.2025.104687
- Ward-Cherrier, B., Conradt, J., Catalano, M. G., Bianchi, M., and Lepora, N. F. (2020). “A miniaturized neuromorphic tactile sensor integrated with an anthropomorphic robot hand,” in *Proceedings of the 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 9883–9889. doi: 10.1109/IROS45743.2020.9341391