OPTICA

Self-calibrated neuromorphic hyperspectral derivative imaging

RONGZHOU CHEN, CHUTIAN WANG, YUXING LI, YUQING CAO, SHUO ZHU, AND EDMUND Y. LAM*

Department of Electrical and Electronic Engineering, University of Hong Kong, Pokfulam, Hong Kong SAR, China
**elam@eee.hku.hk*

Received 2 December 2025; revised 26 February 2026; accepted 26 February 2026; published 25 March 2026

Hyperspectral imaging provides rich spectral-spatial contrast, but conventional systems are often limited by cost, acquisition speed, and dynamic range. We present OpenHI, a low-cost event-based hyperspectral imager that uses an event camera during spectral scanning to capture spectral-derivative frames with respect to the wavelength of the scene's log-transmittance, achieving microsecond latency and high dynamic range. To address distortions from non-constant scan velocity, including acceleration-induced temporal shear, we apply a multi-window warping model for compensation. The system also automatically recovers scan timing and performs time-wavelength alignment. We compare reconstructed spectral slices with a reference hyperspectral camera under low-light conditions and demonstrate an open, low-cost hardware–firmware–software pipeline that makes neuromorphic hyperspectral imaging more accessible.

https://doi.org/10.1364/OPTICA.585766

Scanning hyperspectral imagers deliver high spectral resolution at the cost of mechanical complexity and strict synchronization between motion and detection elements [1]. Conventional systems integrate encoder signals, tachometers, or trigger wiring to maintain spatial-spectral registration, increasing system cost and reducing robustness in long-term or field deployments. Event cameras offer a complementary sensing modality: they asynchronously record logarithmic intensity changes with microsecond latency and high dynamic range, capturing only the temporal variations induced by a scanning illumination pattern [2]. Its high dynamic range and microsecond-level latency enable reliable detection of intensity changes even under dim or rapidly varying illumination, and in strong ambient light, all without requiring exposure control. Prior event-driven hyperspectral work has typically relied on actively modulated illumination or precisely calibrated scanning optics to ensure temporal alignment [3].

Here, we introduce OpenHI, a self-calibrated hyperspectral system that allows non-constant motor motion and restores spectral structure algorithmically. The design emphasizes low complexity and accessible hardware while preserving spectral-derivative fidelity; when sourced from commodity suppliers, the core scanning illumination module can be built for about USD 35 (excluding the event camera and optional $4f$ validation optics), with details in Table S3 of Supplement 1. The system operates in transmission across the visible band (420–700 nm), reducing surface-scatter artifacts and enabling a tractable model that links event timing to wavelength. The key contributions of OpenHI are: (i) a self-segmentation method that separates forward and backward scans directly from the event activity trace without external triggers; (ii) a multi-window time warp model with soft membership functions that minimizes within-bin variance to correct distortions from scanning, including acceleration-induced temporal shear, extending affine warps from event-motion studies [4,5]; and (iii) a reconstruction pipeline that integrates polarity-balanced events into spectral estimates, benchmarked against a frame-based hyperspectral reference.

Figure 1 summarizes the neuromorphic hyperspectral microscope. As illustrated in Fig. 1(a), the dispersed illumination is mechanically swept across the specimen, such that each spatial location experiences a time-varying wavelength. Although our compensation framework applies to motion-induced distortions along arbitrary directions, for clarity, we model the system along the scan axis x and denote the incident wavelength by [6,7]

$$\lambda(x, t) \approx d\left(\frac{x}{z_2} - \frac{\xi(t)}{z_1}\right), \tag{1}$$

where d is the grating period, z_2 and z_1 respectively denote the distances from the grating to the sensor and illumination relay, and $\xi(t)$ is the motor displacement (see Supplement 1 for details). The transmitted intensity can be written as $I_{x,y}(t) = I_0^{(x,y)}[\lambda(x, t)]$, leading to the temporal gradient:

$$\frac{\partial I_{x,y}}{\partial t} = \frac{\partial I_0^{(x,y)}}{\partial \lambda}\frac{\partial \lambda(x, t)}{\partial t}; \quad \frac{\partial I}{\partial t}\Delta t \approx \Delta I, \tag{2}$$

which drives event emission when the logarithmic intensity crosses the sensor threshold [2,8], producing the asynchronous event cloud visualized in Fig. 1(b). Because events encode changes in log-intensity, we reconstruct hyperspectral derivative images from binned event accumulations; recovering absolute intensity would require additional reference information (see Supplement 1 for notation and details) [2,3]. This pipeline enables the label-free spectral imaging of diverse biological specimens, such as planaria and organoids [Figs. 1(c) and 1(d)].

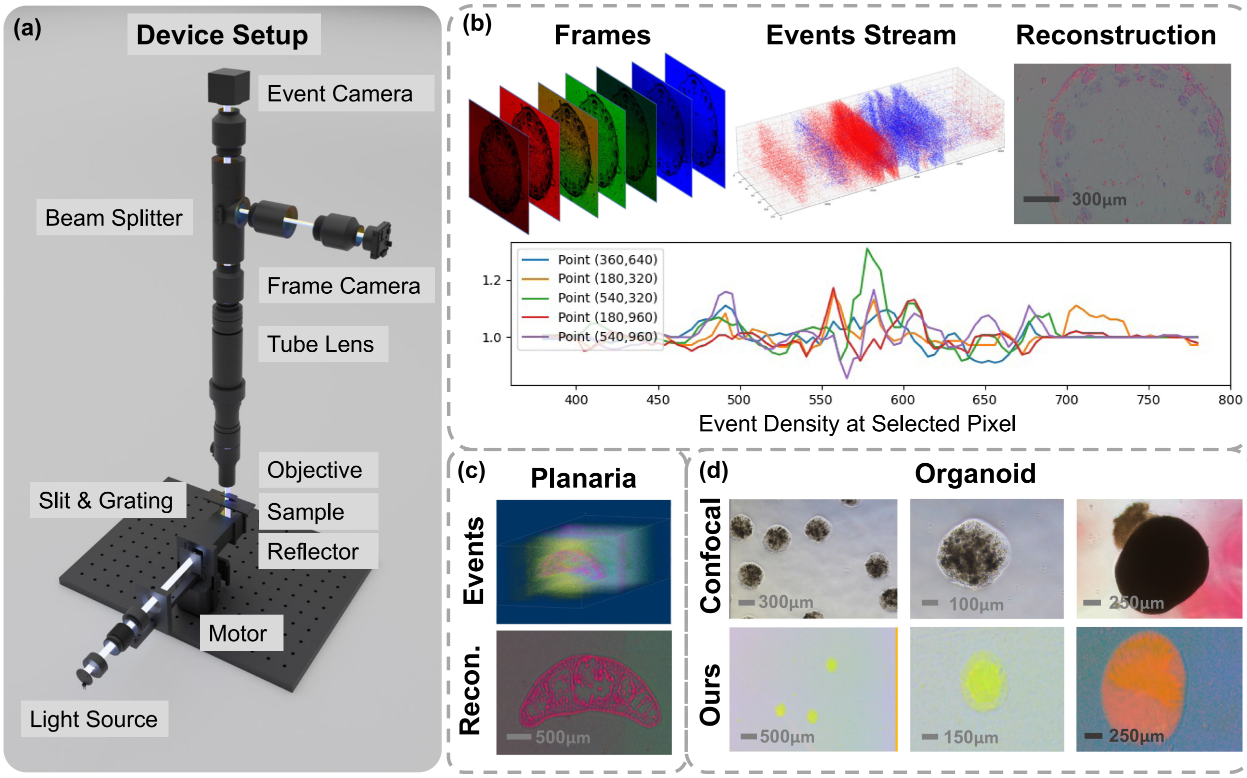

Fig. 1. System overview and modular integration. (a) A dispersed illumination path acts as a drop-in module before the sample, while a non-intrusive detection add-on uses a 4 f relay to feed the event and frame sensor, with a beam splitter preserving the microscope's original imaging path. (b) Frames, 3D event cloud, and representative reconstructed hyperspectral slices (mapped to RGB using a CIE-based conversion), and example event-density traces from binned frames. (c) Event recordings of planaria and their CIE-mapped RGB reconstructions. (d) Organoids imaged with confocal microscopy at 4× (top row) and with the 1× transmission event-based mode (bottom row).

We bin event timestamps into a 1 ms activity trace to estimate scan timing. Auto-correlation of this trace and correlation with the time-reversed signal (auto-convolution) [9] yield the scan period and the start, turnaround, and end boundaries without external timing signals (Fig. 2). We use these peaks to segment forward and backward passes for reconstruction, and details and robustness checks are provided in Supplement 1.

Even after accurate segmentation, residual non-constant motor motion introduces spatially varying temporal shear beyond the nominal scan geometry in Eq. (1) [5,6]. To address this, we model each scan as a sequence of M temporal windows in which the effective scan speed is approximately stable; windows are blended via soft membership weights $w_i(x, y, t)$ to avoid discontinuities. Given a set of warp parameters $\theta = \{(\tilde{a}_i, \tilde{b}_i)\}_{i=1}^{M}$, where $\tilde{a}_i$ and $\tilde{b}_i$ denote the lateral and vertical time-per-pixel slopes in window i, the raw events are mapped from their original timestamps t to compensated timestamps:

$$t'(x, y, t; \theta) = t - \sum_{i=1}^{M} w_i(x, y, t)\left(\tilde{a}_i x + \tilde{b}_i y\right), \tag{3}$$

where $w_i(x, y, t)$ denotes the soft assignment weight of the event to temporal window i, with $\sum_i w_i = 1$. The compensated events are then accumulated into frames $\mathsf{E}_k(x, y; \theta)$ for each temporal bin k. We optimize θ by minimizing

$$\mathcal{L}(\theta) = \underbrace{\frac{1}{HW}\sum_{k}\sum_{x,y}\left(\mathsf{E}_k(x, y; \theta) - \mu_k\right)^2}_{\mathcal{L}_{\mathrm{var}}} + \gamma_{\mathrm{sm}} \underbrace{\sum_{i=1}^{M-1}\left[(\tilde{a}_{i+1} - \tilde{a}_i)^2 + (\tilde{b}_{i+1} - \tilde{b}_i)^2\right]}_{\mathcal{L}_{\mathrm{sm}}}, \tag{4}$$

where H and W are the image height and width, μ_k is the spatial mean of frame k, and γ_{sm} controls the strength of the smoothness penalty that encourages neighboring windows to have similar slopes. The coefficients $\tilde{a}_i$ and $\tilde{b}_i$ correct for effective scanning speed variations. The data fidelity term $\mathcal{L}_{\mathrm{var}}$ favors sharp, spatially coherent accumulations, while the smoothness penalty $\mathcal{L}_{\mathrm{sm}}$ prevents overfitting by enforcing slowly varying slopes across adjacent windows. Minimizing $\mathcal{L}(\theta)$ by gradient descent aligns the temporal windows, suppresses acceleration-induced shear, and yields sharper spectral reconstructions.

Figure 3 summarizes the method for a cross-section of a dicot stem. Figures 3(a) and 3(b) show 3D event clouds before and after the learned time-warp, highlighting the removal of spatial–temporal shear from motor motion. In each panel, a dashed box marks the spatial extent of the 50 ms accumulation window used for comparison. Figures 3(c) and 3(d) present the corresponding event accumulations before and after compensation over the same window. These panels show that the multi-window model corrects motion-induced distortions across the view.

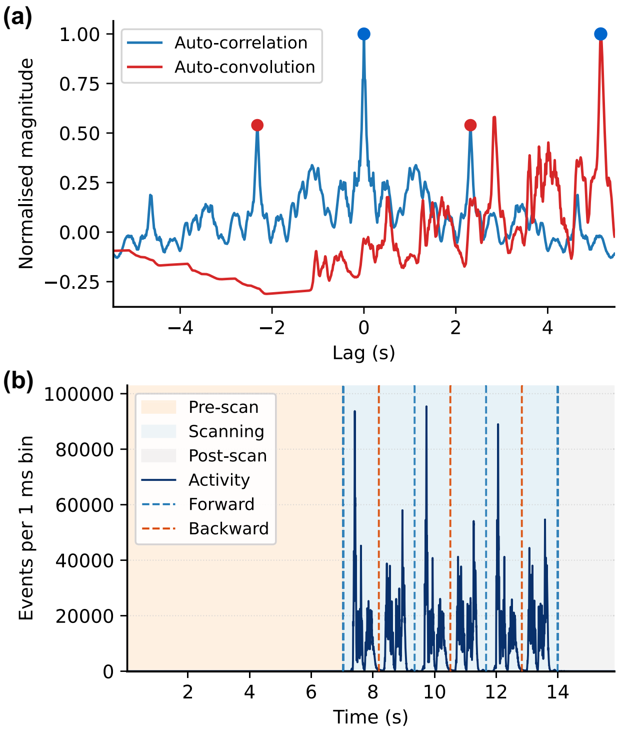

Fig. 2. Self-calibrated event detection and scan segmentation. (a) Auto-correlation and auto-convolution of the event activity signal; the red and blue markers indicate the peaks used to determine the period, start, and end of scanning. (b) Event counts over the full recording with pre- and post-scan shading and alternating forward and backward boundaries.

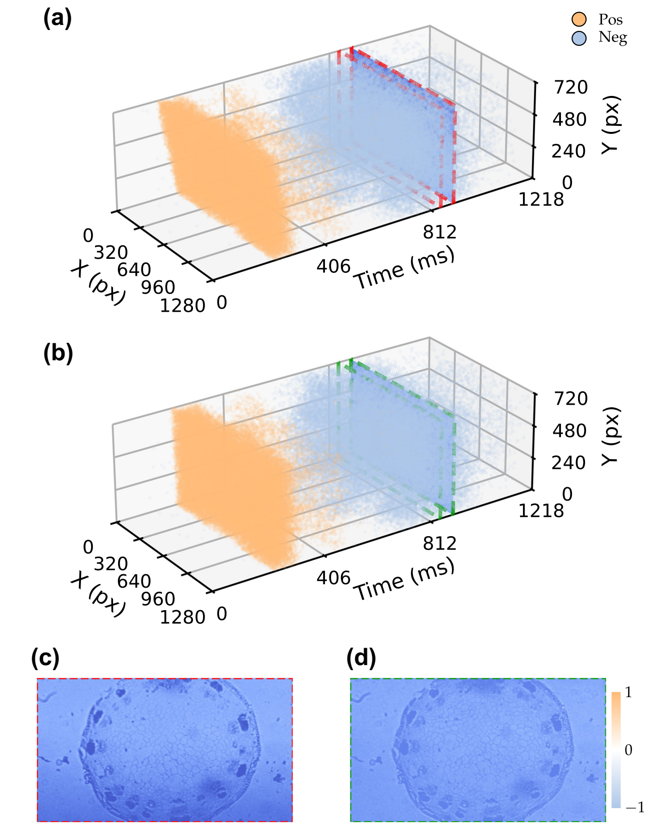

Fig. 3. Multi-window scanning compensation. (a, b) 3D event clouds before/after compensation; dashed outlines mark the 50 ms accumulation window used for comparison. (c, d) Corresponding normalized event accumulations over the same 50 ms window.

With temporal distortions corrected, the compensated event stream removes the spatial dependence of wavelength arrival, allowing λ to be modeled as a function of t' alone. We compute a one-dimensional background trace $b(t')$ from polarity-weighted compensated events in 5 ms bins; robustness across a blank slide and a representative sample region is shown in Supplement 1 (Fig. S6). After normalizing $b(t')$ between pre/post-scan plateaus, we identify the active scan interval $[t_0, t_1]$ from its rising and falling edges. We use this illumination reference to identify the effective start and end of the scanned spectral band and to define a consistent wavelength range for binning. Using reference spectra with known wavelength endpoints, we fit an affine mapping $\lambda(t) = \alpha t + \beta$ over this interval. The normalized background trace is then reparameterized from time to wavelength using this mapping and matched to the mean ground-truth spectrum via a least-squares fit for amplitude. This calibrated time-wavelength mapping underpins the spectral sequence in Fig. 4, and a confidence-based spectral-resolution discussion for the derivative modality is provided in Supplement 1.

After compensation, events are accumulated along t' into quasi-static frames $\mathsf{E}_k(x, y)$, with each cumulative window advanced by a 5 ms temporal step, and polarity-weighted counts integrated to recover relative spectra. With the temporal alignment and wavelength calibration in place, we reconstruct the spectral sequence. Figure 4 illustrates the reconstructed slices alongside measurements from a conventional hyperspectral camera. In Fig. 4, the event result is shown as accumulated events with background-subtracted to remove stationary background activity, whereas the frame reference is shown as linear-intensity finite differences. The top two rows show 50 ms event accumulations before and after multi-window compensation; the third row presents the corresponding 20 nm finite-difference gradients derived from the reference hyperspectral stack. As highlighted by the dashed box, the event-derived slices reveal clearer wavelength-dependent spatial features than the reference gradients. We attribute this difference to the fact that computing finite differences from low-light frame-based measurements amplifies read noise and fixed-pattern artifacts (e.g., stripes), whereas the event sensor directly detects relative intensity changes, effectively bypassing this noise-amplification step [10]. Under identical illumination, the event-driven system completes a spectral scan in about 585 ms, compared with roughly 300 s for a conventional frame-based hyperspectral camera. The resulting data volume is also substantially smaller (about 18.5 MB versus 138 MB), which streamlines transfer and storage for portable use. The optical module is built from simple components—an LED, diffraction grating, and low-cost motor—with structural parts produced by 3D printing; costs are reported in Table S3 of Supplement 1 (core module) and listed separately from the optional $4f$ validation optics. In addition to the drastic reduction in acquisition time and cost, the asynchronous sensing architecture offers high dynamic range and resilience to illumination fluctuations, maintaining reliable performance even in low-light environments where the frame-based camera exhibits noise saturation, striping artifacts, and motion blur.

The results above show that event-driven hyperspectral derivative acquisition is possible without encoder synchronization or

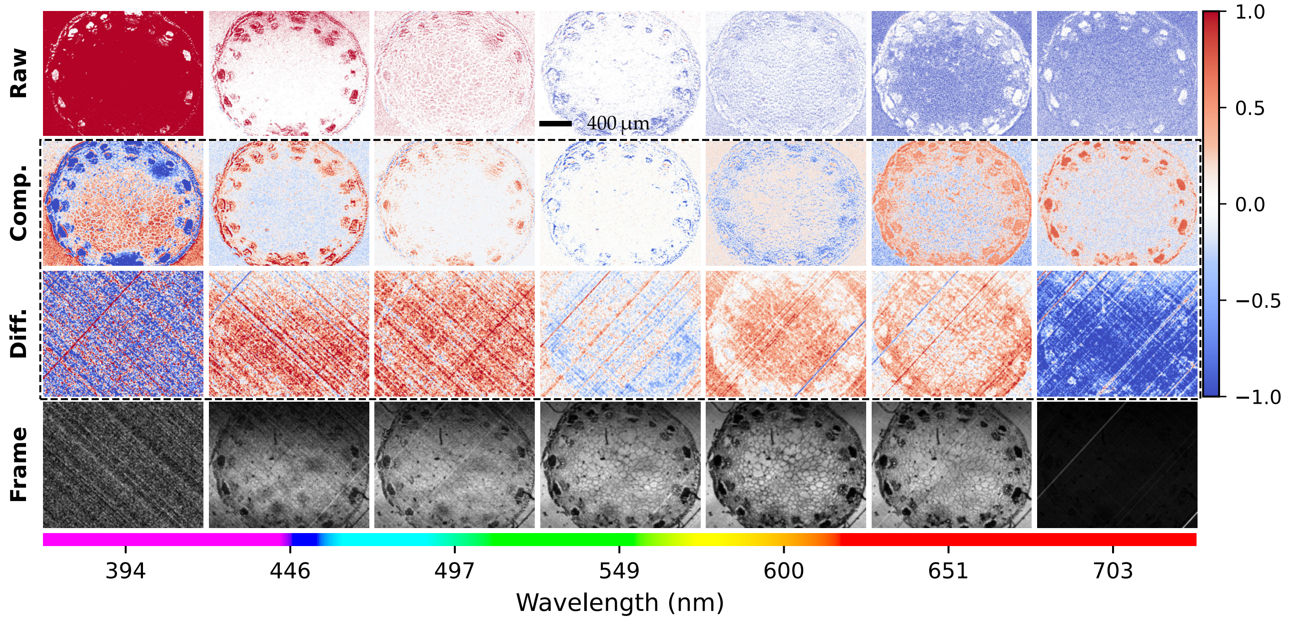

Fig. 4. Spectral reconstruction results (400–700 nm). Columns show seven 50 ms bins spanning 394–703 nm. Events are binned by corrected time t' and mapped to wavelength via $\lambda(t) = \alpha t + \beta$ (Supplement 1). Rows (top to bottom): raw event accumulations; compensated accumulations with background-subtracted event activity; 20 nm finite-difference gradients from the reference hyperspectral camera; and frame-based reference exposures.

precision motion control. By inferring scan timing directly from the event stream and correcting velocity-induced distortions with a variance-minimizing multi-window time warp, OpenHI recovers wavelength-dependent derivative structure from compact, asynchronous measurements. This approach adapts event-based motion-compensation concepts to the spectral domain and enables stable, gradient-based spectral representations when illumination is limited. A simple optical layout, open hardware–firmware–software design, and algorithmic self-calibration suggest that neuromorphic hyperspectral derivative imaging can be implemented in lightweight, low-cost form factors. These characteristics point to accessible systems suitable for portable or field-deployable platforms and highlight the potential of event-driven sensing as a foundation for future compact spectral imagers [11–13], Code 1, Ref. [14]; Code 2, Ref. [15]; Code 3, Ref. [16]; and Code 4, Ref. [17].

Funding. University Research Committee, University of Hong Kong (URC 2402101794); Research Grants Council of the Hong Kong Special Administrative Region, China (Theme-based Research Scheme, T45-701/22-R).

Acknowledgment. We thank Davit Khutsishvili, Dr. Zitian Wang, Lijuan Zheng, Lingxiao Yang, and Prof. Shaohua Ma (Tsinghua Shenzhen International Graduate School, Tsinghua University) for supplying organoids, providing the confocal images, and other support. We further acknowledge Dr. Yanmin Zhu and Yunfei Tian for helpful assistance and discussions.

Disclosures. The authors declare no conflicts of interest.

Data availability. All design assets (3D-print files, LED PCB, motor firmware, and acquisition GUI for Prophesee/DAVIS) are provided in Supplement 1 and mirrored at Code 1, Ref. [14]; Code 2, Ref. [15]; Code 3, Ref. [16]; and Code 4, Ref. [17].

Supplemental document. See Supplement 1 for supporting content.

REFERENCES

1. N. Hagen and M. W. Kudenov, Opt. Eng. **52**, 090901 (2013).
2. G. Gallego, T. Delbrück, G. Orchard, *et al.*, IEEE Trans. Pattern Anal. Mach. Intell. **44**, 154 (2022).
3. R. Liu, Y. Xiao, Z. Wang, *et al.*, Opt. Lett. **48**, 3387 (2023).
4. C. Wang, S. Zhu, P. Zhang, *et al.*, Laser Photonics Rev. **19**, 2400647 (2025).
5. G. Gallego, H. Rebecq, and D. Scaramuzza, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, 2018), pp. 3867–3876.
6. C.-K. Yeh and V. J. Tsai, in *2011 IEEE International Geoscience and Remote Sensing Symposium* (IEEE, 2011), pp. 2881–2883.
7. G. Zhou, K. K. Cheo, Y. Du, *et al.*, Opt. Lett. **34**, 764 (2009).
8. B. Yu, J. Liang, Z. Wang, *et al.*, in *Proceedings of the Computer Vision and Pattern Recognition Conference* (IEEE, 2025), pp. 929–939.
9. R. Gorenflo and B. Hofmann, Inverse Probl. **10**, 353 (1994).
10. R. Cao, D. Galor, A. Kohli, *et al.*, Optica **12**, 46 (2025).
11. S. Zhu, Q. Yin, C. Wang, *et al.*, Adv. Sci. **12**, e10338 (2025).
12. Z. Ge, C. Wang, J. Huang, *et al.*, Opt. Lett. **50**, 1496 (2025).
13. L. Bian, Z. Wang, Y. Zhang, *et al.*, Nature **635**, 73 (2024).
14. R. Chen, C. Wang, Y. Li, *et al.*, "Self-calibrated neuromorphic hyperspectral derivative imaging," github, 2025, https://github.com/lachlanchen/OpenHI.
15. R. Chen, C. Wang, Y. Li, *et al.*, "3D mechanical design files for the OpenHI core neuromorphic hyperspectral module," figshare, 2025, https://doi.org/10.6084/m9.figshare.31289644.
16. R. Chen, C. Wang, Y. Li, *et al.*, "3D CAD models for the optional 4f optical relay module," figshare, 2025, https://doi.org/10.6084/m9.figshare.31289650.
17. R. Chen, C. Wang, Y. Li, *et al.*, "PCB design files for the LED illumination module (LED2835)," figshare, 2025,, https://doi.org/10.6084/m9.figshare.31289656.