

UC Berkeley

UC Berkeley Previously Published Works

Title

Spectral kernel machines with electrically tunable photodetectors

Permalink

<https://escholarship.org/uc/item/48c1820m>

Journal

Science, 390(6776)

ISSN

0036-8075

Authors

Zhang, Dehui

Li, Yuhang

Geng, Jamie

et al.

Publication Date

2025-11-27

DOI

10.1126/science.ady6571

Peer reviewed

Title: Spectral kernel machines with electrically tunable photodetectors

Authors: *Dehui Zhang^{1,2,3}, Yuhang Li⁴, Jamie Geng^{1,2,3}, Hyong Min Kim^{1,2,3}, Marco Ma¹, Shifan Wang^{1,2,5}, Inha Kim^{1,2,3}, Theodorus Jonathan Wijaya^{1,2,3,6}, Naoki Higashitarumizu^{1,2,3}, I. K. M. Reaz Rahman^{1,2,3}, Dorottya Urmossy¹, James Bullock⁵, Aydogan Ozcan^{4,7,*}, and Ali Javey^{1,2,3,8,*}*

Affiliations:

¹ Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA, 94720, USA.

² Materials Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA.

³ Berkeley Sensor & Actuator Center, University of California, Berkeley, CA, 94720, USA.

⁴ Department of Electrical and Computer Engineering, University of California, Los Angeles, CA, 90095, USA.

⁵ Department of Electrical and Electronic Engineering, University of Melbourne, Melbourne, Victoria, 3010, Australia.

⁶ Department of Electrical Engineering and Information Systems, The University of Tokyo, Tokyo, 113-8656, Japan

⁷ California NanoSystems Institute (CNSI), University of California, Los Angeles, CA, 90095, USA.

⁸ Kavli Energy NanoScience Institute, University of California, Berkeley, CA 94720, USA.

* Address correspondence to: ozcan@ucla.edu, A.O., ajavey@berkeley.edu, A. J.

Abstract: Spectral machine vision collects spectral and spatial information as 3D hypercubes and digitally processes them, which causes a data bottleneck, limiting power efficiency, frame rate, and spectral-spatial resolution. This work introduces spectral kernel machines (SKM) to overcome these bottlenecks. SKM directly compresses spectral analysis through the output photocurrent and learns from example objects to identify/classify new samples in a ‘sniff-and-see’ mode. We experimentally demonstrated SKM with electrically tunable bipolar black phosphorous (bP)-MoS photodiodes in the near/mid-infrared band and silicon photoconductors in the visible band, performing versatile intelligent tasks from chemometrics to semiconductor metrology. This architecture consumed significantly lower power and was more than an order of magnitude faster than existing solutions for hyperspectral image analysis, defining an intelligent imaging and sensing paradigm with intriguing possibilities.

Main text:

Spectral imaging technologies produce data volumes that are orders of magnitude larger than standard color (red/green/blue) image channels, forcing technologies to make trade-offs between spatial/spectral resolution, frame rate, and power efficiency (¹⁻⁵). This can be undesirable for real-time applications that demand fast, continuous interactions with the surrounding environment, such as precision agriculture (⁶), rescue missions (⁷), and autonomous driving (⁸) scenarios. Furthermore, spectroscopy tools should manage unpredictable situations in field deployments, such as the absence of fully known ground truth or target spectra, shifting targets, and ever-changing environmental conditions (⁹).

In-sensor computation could learn to accommodate various sources of changes and compute the desired inference or analysis, mitigating data-storage/transmission/processing-related

barriers. Recent neuromorphic photodetectors demonstrated spatial pattern identification (¹⁰⁻¹²) and intensity adaptation (¹³⁻¹⁷). For spectral vision, earlier attempts demonstrated certain capabilities, such as color adaptation or addition/subtraction with two or three visible colors/bands (¹⁸⁻²¹), but often sacrificed spatial resolution in favor of spectral resolution (Table S1).

Here, we show that an intelligent photodetector design can learn to perform machine learning analysis over complex, visible to mid-infrared (MIR) incident spectra, with the readout photocurrent directly providing the final inference. This process mathematically resembles the kernel machine algorithms typically used for digital machine learning, making each photodetector a spectral kernel machine. It can ‘sniff-and-see’ conceptually inspired by retriever dogs, learning from examples to recognize spectral features within a complex scene. We experimentally demonstrated different SKMs to perform wafer metrology and feature identification with a bandwidth-limited inference time of 2 μ s in the visible band. For the MIR band, an electrically tunable bipolar photodiode made of bP-MoS₂ heterostructure was operated at room temperature to identify chemicals and quantify mixture concentrations indistinguishable in the visible bands. The demonstrated in-device intelligence creates new opportunities for spectral machine vision applications that demand low power, high speed, and high resolution, which are especially attractive for implementations in, e.g., mobile devices, robotics, and satellites.

Spectral kernel machines

We show that the spectral feature $P(\lambda)$, representing the wavelength-dependent reflectance/transmittance of a static test object, can be linearly superposed within a photodetector to enable separation between target classes of interest in a transformed domain (Fig. 1A). To achieve

this separation, we first encoded $P(\lambda)$ into a time-varying optical intensity signal $P(\lambda) \cdot F(\lambda, t)$, where $F(\lambda, t)$ denotes the encoding function. This encoding can be achieved using modulated illumination (e.g., fast flashing LEDs with different colors) or modulated transmission (e.g., a rotating color filter disk in front of the detector). Simultaneously, the photodetector's responsivity R was controlled by a voltage $V_{ctr}(t)$ (Figs. 1B, C) and synchronized with $F(\lambda, t)$. Consequently, the resulting photocurrent $\overline{I_{ph}}$, time-averaged by the system's bandwidth, can be written as:

$$\overline{I_{ph}} = \iint P(\lambda) \cdot F(\lambda, t) \cdot R(\lambda, V_{ctr}(t)) \cdot \Gamma(t) \cdot d\lambda \cdot dt \quad (1)$$

where $\Gamma(t)$ represents the temporal averaging function due to either the detector's intrinsic RC delay or lower-speed readout circuits; see Fig. 1D.

We define a task-specific spectral function $Q(\lambda)$, trainable by tuning $V_{ctr}(t)$, as:

$$Q(\lambda) = \int F(\lambda, t) \cdot R(\lambda, V_{ctr}(t)) \cdot \Gamma(t) \cdot dt \quad (2)$$

Substituting this equation into Eq. 1, the total current becomes:

$$\overline{I_{ph}} = \int P(\lambda) \cdot Q(\lambda) \cdot d\lambda \quad (3)$$

As detailed in the Supplementary Materials, $Q(\lambda)$ can be reconfigured via $V_{ctr}(t)$ to approximate various spectral functions. We can fit $Q(\lambda)$ to a weighted superposition of training spectra: $b(\lambda) = \sum_{i=1}^N w_i z_i S_i(\lambda)$, where $\{S_i(\lambda)\}$ represents the known spectral features of

$i = 1, \dots, N$ training example objects, $\{w_i\}$ are the trainable weights, and $z_i \in \{-1, +1\}$ indicates whether the training example is negative or positive. Accordingly, Eq. 3 can be written as:

$$\overline{I}_{ph} \approx \sum_{i=1}^N w_i z_i k(P, S_i) \quad (4)$$

where $k(P, S_i)$ defines an inner product in the spectral hyperspace:

$$k(P, S_i) = \int P(\lambda) \bullet S_i(\lambda) \bullet d\lambda \quad (5)$$

This linear kernel measures the similarity between the spectrum of an object under test and the known spectral features, $S_i(\lambda)$. $k(P, S_i)$ can also be made nonlinear with higher-order photo physics⁽²²⁻²⁴⁾ or readout nonlinearities, although not considered in this work.

Eq. 4 connects the measured current $\overline{I}_{ph}$ with $P(\lambda)$ and $\{S_i(\lambda)\}$, analogous to the kernel machines^(25, 26) commonly used for classification and regression tasks. In a kernel machine for binary classification, the class of an unknown sample Φ is predicted by:

$$z = \text{sign} \left(\sum_{i=1}^N w_i z_i k(\Phi, \Phi_i) \right) \quad (6)$$

where $\{\Phi_{i=1, \dots, N}\}$ is a set of N training examples. The kernel function $k(\Phi, \Phi_i)$ measures the similarity of Φ to Φ_i . $\{w_i\}$ are trainable weights, and $z \in \{-1, +1\}$ indicates whether the example is negative or positive, sharing the same formalism as in Eq. 4.

Therefore, our SKM device, by its design, implements a kernel machine-based classifier. The binary label of the test object is given by the sign of $\overline{I_{ph}}$:

$$z = \underset{(7)}{\text{sign}}(\overline{I_{ph}})$$

Geometrically, each $P(\lambda)$ can be interpreted as a point in a spectral hyperspace, constructed by $\{P(\lambda_k)\}$ (fig. S1). $Q(\lambda)$ defines a decision boundary - a hyperplane that separates positive and negative examples. The analog computation of $\overline{I_{ph}}$ is equivalent to projecting $P(\lambda)$ along a direction perpendicular to this maximum-margin hyperplane and making classification decisions based on the sign of the photocurrent.

Such an operation principle demands photodetectors with electrically tunable responsivity and photocurrent that can flow in opposite directions, enabling $Q(\lambda)$ to take both positive and negative values for general computation. The dark current must also be much smaller than the photocurrent to avoid baseline drift and poor signal-to-noise ratio (SNR). We solved these challenges with hardware-software co-design, as detailed in the following sections and Supplementary Materials.

Training and testing of visible-band SKMs

The traditional backpropagation-based training methods cannot be applied to SKM because the gradients of computational nodes are not accessible. Moreover, the drifting and noise in the analog hardware are hard to model accurately at the software level. Hence, we need a gradient-blind, in-situ training algorithm based on real-time feedback loop between the measurements and control software (^{27, 28}) to naturally accommodate these nonidealities. Through a ‘sniff-and-see’ operation

(Fig. 2A, C), we enabled SKMs to learn from examples and then perform multiple machine vision tasks as detailed below.

We first used the SKM framework to demonstrate image segmentation with a color-printed photo of a golden pheasant (*Chrysolophus pictus*) based on its spectra (fig. S2). Six fast-flashing visible LEDs with different spectra (Fig. 2D) illuminated the object. The reflected light was detected by a photoconductor with $V_{ctr}(t)$ as the source-drain bias. Each sniff optically exposed the device to 80 random positive and negative pixels from the training images (Fig. 2A), with 50 trial $V_{ctr}(t)$ sent to the SKM device. A computer collected $\overline{I_{ph}}$ from the SKM for each $V_{ctr}(t)$ and evaluated the prediction accuracy. A genetic algorithm (²⁹) then used better-performing waveforms to generate new ones sent to the device for the next sniff (fig. S3). After three sniff cycles, the optimized $V_{ctr}(t)$ waveforms for highlighting the regions of the bird or leaves (Fig. 2B) were sent to the SKM. This process led to weighted task-specific spectral functions $Q(\lambda)$, highlighting the target bands and suppressing everything else (Figs. 2E, F). Figures 2G-J show the identification results for leaf/bird regions, demonstrating an intelligent spectral selectivity. The testing pixels belonged to an image never seen by the SKM before, demonstrating the success of our blind testing. In comparison, other passive filtering and thresholding methods showed no reconfigurability and poorer accuracy (see fig. S4).

To further demonstrate how the system can also benefit from digital back-end processing, Fig. 2K integrated a lightweight and shallow convolutional neural network (CNN) with only 21.8k trainable parameters as a digital post-processor of the SKM system. Once the CNN was trained (see Materials and Methods), the segmentation accuracy increased from 96.53% (SKM alone) to 98.57% by harnessing both the spectral and spatial features of the object.

As another test-bed, we used the SKM framework to identify different SiO₂ thickness values on silicon wafers based on their thin film reflection spectra (fig. S6A). Despite the limited number of spectral bands available, the task-specific $Q(\lambda)$ automatically approached the desired algorithmic function $b(\lambda)$ for classification (fig. S6B) after training (see Table S2 for details), and identified the positive samples with 100% accuracy over 21 oxide thickness values tested on new unseen wafers (100-300 nm, 10 nm step) (fig. S6D). We further explored the accuracy of SKM for inspecting small thickness variations. After training on some example wafers (Fig. 3A), the SKM was blindly tested on new samples and was able to identify 5 Å thickness difference at 32 sampling points on two wafers never seen by the SKM in the training process, with only two misclassified points (Fig. 3B), demonstrating accurate measurement of sub-nanometer thickness variations in this blind testing. In another test, when V_{ctr} was broadcast for a single cycle for $\sim 1.6 \mu\text{s}$, the net charges built up to positive for the 140-nm SiO₂ sample (trained as positive, blindly tested on a new wafer), and fell back at $\sim 1.9 \mu\text{s}$, but were around zero for 50-nm SiO₂ (fig. S7A). The rise/fall times of the recognition process approached the detection speed (fig. S7B), because $V_{ctr}(t)$ smartly manipulated the transient photo charge dynamics even before full photocurrent extraction with complete rise and fall slopes (figs. S7C, D). We also trained another SKM to reveal the oxide thickness directly as the readout current (1 nA/nm, Fig. 3C), with the training and test settings reported in Table S2. The mean error and the standard deviation values of the SKM inference on new test samples, never seen before, were found to be 0.3 nm and 1.16 nm, respectively, experimentally demonstrating its accuracy and precision under blind testing.

As another test-bed, we demonstrated that our SKM framework can also distinguish natural and artificial leaves. Figure 3D shows natural leaves from 7 different plants and artificial

polymer/fabric leaves from 4 different manufacturers. Despite their visible resemblance, the artificial leaves showed spectral fingerprints different from natural leaves (see e.g., fig. S8A). The trained SKM was blindly tested on 18 new natural leaves and 13 new artificial leaves, never seen before, (Fig. 3E), achieving a classification accuracy of 93.5%. The SKM accommodated the species-to-species and leaf-to-leaf variations in the same class, demonstrating its generalization capabilities.

As a related test-bed, we also used our framework for the classification of leaf hydration levels, which served as a key indicator determining irrigation needs, potentially facilitating sustainable usage of water in agriculture. For this task, we retrained the SKM hardware to classify which leaves of the tree *Styphnolobium japonicum* were of low hydration levels based on their spectral information (see e.g., fig. S8B); the ground truth class for each leaf was measured using standard water concentration calibration procedures as detailed in the Materials and Methods section. The $\overline{I_{ph}}$ values from 24 test leaves (never seen before) are plotted in Fig. 3F with respect to the relative water concentration of each leaf. These blind testing results demonstrated a binary classification accuracy of 95.8%. Table S2 summarizes the training details for all the experiments.

MIR SKM with gate-tunable bP-MoS₂ photodiodes

SKM devices working with MIR spectra are especially attractive for chemometrics and biosensing due to the distinct spectral signatures of various materials available in this band (³⁰⁻³²). Photoconductive devices are not well suited for this purpose, as narrow bandgap semiconductors are required for MIR detection, resulting in high dark current (³³⁻³⁷), drift and high noise at room temperature (³⁸⁻⁴⁰). We addressed the problem with a gate-tunable bipolar bP-MoS₂ photodiode. bP-MoS₂ photodiodes showed room-temperature MIR photovoltaic response due to the high MIR

quantum efficiency of bP and matched band offset between bP and MoS₂ for photo charge separation (⁴¹).

Figure 4A shows the bP-MoS₂ device structure and operation. We applied V_{ctr} as a gate bias and operated the diode at zero source-drain bias (i.e., short-circuit photocurrent mode) to minimize the dark current. MoS₂ screened the overlapping bP from the gate fields. Consequently, a lateral homojunction was formed in bP, with the built-in voltage polarity and magnitude dependent on the applied V_{ctr} . Figure 4B depicts the simplified band diagrams within the lateral bP homojunction. At $V_{ctr} = 0$ V, there is a p⁺-p⁻ junction in bP. The photo-generated holes move to the left contact, and MoS₂ collects the excess electrons. Therefore, the photocurrent flows from MoS₂ to bP, denoted as $I_{ph} < 0$. $V_{ctr} = 0$ V creates an n-p⁻ junction with $I_{ph} > 0$. The measured photocurrent polarity as a function of gate voltage (Fig. 4C) was consistent with the proposed band diagrams. The device's specific detectivity was measured with the gate grounded ($V_{ctr} = 0$ V), which was above 10⁸ Jones at the target spectral range of 2-3.6 μ m (Fig. 4D) measured by Fourier transform infrared spectroscopy (FTIR), providing sufficient photocurrent for the in-sensor computation. More detailed descriptions of the working principles, device simulations, and characterization of the bipolar gated diode are provided in fig. S9.

MIR chemometrics

As a proof of concept, we used our MIR SKM device to analyze four different chemicals: dimethyl sulfoxide (C₂H₆OS), perfluorooctane (C₈F₁₈), glycerin, and silicone oil. These chemicals are visibly transparent liquids (Fig. 4E), but show different characteristic absorption peaks at 2.5-4.0 μ m wavelengths (fig. S10A). Their transmission spectra were encoded by a rotating disk with

six spectral filters showing nontrivial features at the same band (figs. S10B, C). Overlapping these chemicals and filters with the detector, we observed distinct photocurrent features (fig. S11), allowing chemometric classification after two cycles of training (with $N = 2$ in Eq. 6; also see Table S2). Figure 4F shows photocurrents highlighting C_8F_{18} and suppressing C_2H_6OS . With C_8F_{18} in the test holder (red curve), the current rose with illumination from a black body IR source (un-shaded area in the plot). With C_2H_6OS (blue), we observed $\overline{I_{ph}} < 0$ with illumination. Figure 4G shows the full binary classification results, with the vertical labels representing the suppressed chemicals and the horizontal labels representing the highlighted chemicals. The trend is consistent for all trained SKM outputs in the chart, demonstrating our success in binary chemical identification. Figure S12 further highlights additional experiments classifying multiple chemicals.

In addition to the binary classification of materials, the same SKM device was experimentally trained for quantitative inference to reveal the chemical concentration using the output photocurrent directly. Figure 4H shows the SKM device detecting concentrations of water mixed in acetone compared to the ground truth. The SKM predicted the unknown concentrations using the photocurrent magnitude with a slope of 1 pA per percentage of water, experimentally achieving a root-mean-square error of 2.49% in volume ratio. The MIR chemometrics tests were performed following the standard sample preparation protocols, suitable for dynamic monitoring of chemicals in controlled environments. Additional efforts, using e.g., reinforcement learning, could be used in future applications to cover uncontrolled testing environments further.

Scaling-up for hyperspectral machine vision

To explore the potential feasibility of the presented framework in more complex environments, we numerically simulated SKM’s performance with the Salinas Scene hyperspectral

dataset of landscapes with vegetables, vineyards, and bare soil, separated into 16 classes (⁴²). We modeled an SKM working at a wavelength range of 0.4-2.5 μm (see Materials and Methods) and trained it with a genetic algorithm with 32 genes, corresponding to 32 down-sampled Gaussian bands (see Table S2). Figure 5A shows the binary classification results of the ‘stubble’ region during 100 sniffs, which gradually converged to the ground truth in black and white on the left. Figure S13 and Movie S1 further illustrate training over all 16 classes. The simulation was further extended to city scene recognition (⁴³) and human brain tumor diagnosis (⁴⁴), as shown in Figs. 5B and C. The average sensitivity (Fig. 5D) and specificity (Fig. 5E) values (see Materials and Methods) were ~0.2-3% lower than the digital processing approaches using a linear support vector machine (SVM) (⁴⁵) or a transformer (⁴⁶), but with two orders of magnitude lower power consumption (Fig. 5F) and latency (Fig. 5G) for capturing and analyzing each spectral-spatial hypercube.

SKM’s high energy efficiency and speed arise from the in-sensor spectral computation and data compression. Traditional off-sensor digital analysis causes a significant data bottleneck (Fig. 6A) in the camera series interfaces (CSI), which can contribute to more than 90% of total energy consumption and latency (fig. S17). On the other hand, much less data needs to be transported and converted to the digital domain with in-sensor SKM (Fig. 6B). Consequently, SKM features low energy consumption and latency, as illustrated in Figs. 5F, G, which are crucial for spectral machine vision applications with, e.g., cellphones, satellites, microscopes, and car cameras (colored shades in Figs. 6C, D), but are not readily available with existing commercial hardware (⁴⁷⁻⁵¹) or prior research benchmarks (2, ⁵², ⁵³).

Conclusion

In summary, the readout from a bipolar detector was shown to directly provide the spectral analysis using the sniff and seek learning mode via the spectral kernel machine architecture. The key component of this architecture is an electrically tunable photodetector for which an applied voltage can dynamically control the polarity and magnitude of the photocurrent, optimized for different computational tasks. This characteristic can be achieved with various optoelectronic materials and devices, including a simple photoconductor or a gated photodiode. We demonstrated in-sensor identification of objects, structures, and chemical compositions at visible-to-MIR bands without any post-processing. It offers orders of magnitude higher power efficiency and speed, and large training flexibility to resolve, recognize, and compress high-dimensional spectral-spatial information. Future explorations of on-chip spectral encoding methods and architectures can potentially improve the performance and compactness of SKM. Further modeling and development of devices for the physical realization of nonlinear kernels could also enable more complex SKM algorithms, targeting improved performance in e.g., speed and accuracy. The findings inspire new intelligent sensors with potential applications in precision agriculture, autonomous driving, and chemical sensing and spatial profiling.

References and Notes:

Materials and Methods

1. Synchronization of the spectral encoder and V_{ctr}

The spectral encoder and tunable photodetector were synchronized with an Arduino microcontroller board. For the visible demonstration, a signal generator fed high-frequency square waves to a flip-flops chip (SN74LS74AN) to generate three-bit address signals and drive two analog multiplexers (CD74HC4051). One multiplexer was used to drive the LED arrays to enable fast flashing above the preamplifier bandwidth limit of around 0.6 MHz. The other multiplexer encoded constant voltages from Arduino to a time series of $V_{ctr}(t)$. Hence, the system enabled response speed test up to the bandwidth limit at microseconds despite Arduino's millisecond resolution. For the MIR demonstrations, we used a chopper synchronization signal to locate the filter position and adjusted $V_{ctr}(t)$ directly with the Arduino PWM output.

2. Signal readout

We used both Arduino and Keysight B2901A SMU to measure the readout. For the response time measurement in fig. S7A, we used an oscilloscope to resolve the high-speed response.

3. Light source and spectral encoder settings

For the visible band test, we chose six commercial LEDs. An 852 nm laser diode replaced the yellow LED in the leaf tests (Figs. 3D-F) to better capture the spectral features. For the MIR binary chemical classification, the light source was a black body source encoded with filters, whose transmission spectra are shown in fig. S10. The six filters were attached to a chopper disk that rotates at 1200~3000 rpm. Since the measurement speed was limited by the black body source heating and cooling time in a few seconds, the chopper was sufficiently fast for spectral encoding.

For the acetone-water mixture concentration test, we chose the following filters to distinguish the water absorption: intrinsic silicon, germanium, Thorlabs filters FELH0650, FB750-40, FEL1000, and FEL1350-1. Aluminum irises were applied to filters when necessary, so that the photocurrents through different filters are of the same order of magnitude, which is essential for good SNR. We used glass Petri dishes (30 mm in diameter) to contain 5 mL of the test liquids for Figs. 4F, G, and a large glass beaker in Fig. 4H for finer concentration control.

4. Training the SKMs with genetic algorithm

We applied the pyGAD package in Python for the genetic in-situ training (29). Table S2 shows the training and test settings for all the experiments and simulations in the paper, including the training and test object parameters, number of genes (spectral encoding bands), number of parent genes mating, number of parents kept in each sniff cycle, the mutation rate, and the fitness function.

For experimental binary classification demonstrations in Figs. 2, 3, and 4, we defined the fitness function as:

$$fitness = \frac{e^{-\beta I_{+i}}}{e^{-\beta I_{+i}} + e^{-\beta I_{-i} - \gamma \max(I_{-i}, 0)}} \cdot I_{+i}$$

I_{+i} and I_{-i} are the photocurrent for positive and negative cases. β and γ are parameters assigned based on optimum training performance. The first term rewards larger I_{+i} compared to I_{-i} , and the second term further penalizes the positive I_{-i} values.

In the hyperspectral datasets (Fig. 5 and fig. S13), the F_1 score (54) was used as the fitness function instead of accuracy:

$$F_1 = \frac{2}{recall^{-1} + precision^{-1}}$$

The F_1 score penalizes both the false positives and false negatives, so that it is a better choice for fitness function compared to accuracy in this imbalanced dataset. The down-sampling of spectral bands by a ratio of seven was selected to facilitate training without significant accuracy degradation, which was also verified in fig. S13.

For the SKM integrated with the shallow CNN post-processor (Fig. 2K), $V_{ctr}(t)$ was first trained using a subset of pixels randomly selected from the training images for bird-background segmentation, as previously described and shown in Fig. 2A. After training $V_{ctr}(t)$, we measured the analog SKM output over the entire training images under this optimized control $V_{ctr}(t)$. These outputs (fig. S5, upper panel) were then used to train a compact / shallow CNN, which leverages the spatial features of the samples to refine the classification results and better delineate the bird regions from the background. Specifically, the input to the CNN was the analog output of SKM, i.e., the signal magnitude prior to applying the binary sign function, preserving more information for post-processing. For each pixel, the CNN takes as input a localized 9×9 window centered on that pixel (fig. S5, lower panel), capturing the spatial context needed for improving the classification accuracy. Once trained, the CNN was blindly applied to the SKM output of a new, unseen test image.

5. Leaf hydration level calibration and ground truth measurements

Leaves with different hydration levels were prepared following the standard protocols, with the relative water concentration calibrated by weight (55). Briefly, the leaves of interest for our SKM testing belong to *Styphnolobium japonicum*. The leaves were kept in water overnight to reach the

saturated weights (SW). We recorded the SW values for each leaf and then kept them in the lab environment at 24 °C and 48% relative humidity. They were kept in air for 30 minutes for the high-hydration-level group, or 4 hours for the low-hydration-level group, before measuring the fresh weights (FW). The leaves are then used to train and test the SKM binary classification. After the test, the leaves are dried at 75 °C for 2 hours. After the drying process, further baking at the same temperature did not show observable further decrease of the leaf weight, suggesting full water removal from the leaf tissue. Then the dry weights (DW) were acquired. The relative water content (RWC) values were calculated as:

$$RWC = \frac{FW - DW}{SW - DW} \times 100\%$$

The leaf RWC onset for plant wilting is around 60-70% (56). Consequently, we used 70% as the threshold for the low-hydration-level leaf group. Based on the preparation protocols summarized above, our low-hydration-level leaf samples' RWC values ranged from 59.1% to 67.2%.

6. Fabrication and characterization of the bP-MoS₂ photodiode

For experimental demonstration, we first transferred and patterned exfoliated MoS₂ flakes (10~50 nm thick) on silicon substrates with 50-nm oxide layer. Next, we evaporated metal electrodes (Ni or Ti/Au for different batches) after another set of lithography and liftoff steps. Last, we dry-transferred a thin bP layer (7-17 nm) passivated with a hexagonal boron nitride (hBN) layer to bridge the MoS₂ and the other contact (fig. S9A). The photo response of the bP-MoS₂ was measured in a vacuum chamber to further avoid bP degradation.

A SuperK Fianium supercontinuum laser was used as the light source for Figs. 4C and H (with additional filters discussed earlier) to provide sufficient optical power between visible to $\lambda = 2.3 \mu$

m. Another $\lambda = 1.55 \mu\text{m}$ laser with fast electrical modulation was used as the light source for the measurement of response speed in figs. S9E, F. The MIR chemical classification test required a broader spectral range, so we used a high-power black body source as the light source and filtered higher energy photons to ensure sensitivity to the fingerprint response from the NIR/MIR bands.

7. Photo response simulation for the bP-MoS₂ photodiodes

Gate-tunable photocurrent and band structures of the bP-MoS₂ photodiodes are simulated with the Sentaurus TCAD software. Estimations of $N_A = 5 \times 10^{17} \text{ cm}^{-3}$ in bP, and $N_D = 1 \times 10^{15} \text{ cm}^{-3}$ in MoS₂ were made according to experimental measurements of bP and MoS₂ transistors. The thicknesses of bP and MoS₂ were set at 12 nm and 35 nm, respectively, in accordance with atomic force microscope (AFM) measurements of the real devices.

8. Simulation with hyperspectral datasets

The modeled SKM in Fig. 5A was spectrally encoded with 32 Gaussian bands down sampled from the 224 bands within the dataset. For the down sampling, the center wavelengths of the Gaussians were equally spaced within the wavelength range of 0.4~2.5 μm . The sigma values of the Gaussians were set as half of the closest center wavelength spacings. Consequently, Eq. 1 reduces to:

$$\overline{I_{ph}} = \alpha \sum_{j=1}^{32} V_j \cdot P_j$$

α is a physical pre-factor for the bias voltage V_j , assuming linear dependence of responsivity on the voltage. For simplicity, the dimensionless unit system in the model assigned α as 1. We assumed no pre-knowledge of the input-output mapping of the device, which was true for the physical device, and therefore, the gradient-descent-based optimization was not valid for the ‘sniff-and-see’. The

hyperspectral human brain dataset was similarly down-sampled by a factor of 4 in the spectral domain. The LIB-HSI dataset was not down-sampled.

The selection of the training parameters involved a critical tradeoff between the learning speed and prediction accuracy. To explore this relationship, we analyzed the SKM's performance on the Salinas hyperspectral dataset using different numbers of $V_{ctr}(t)$ waveforms and training samples per cycle (fig. S18). We empirically defined convergence as the point at which the fitness function fails to improve by more than 1×10^{-5} over 30 consecutive training cycles. As expected, more training samples revealed slower convergence, and a greater number of waveforms decreased the iteration number but increased the total number of measurements. With a relatively small number of waveforms, we can still get good performance; for example, an F_I score of 0.73 can be achieved with 850k pixels read from the SKM during 21 training cycles using 100 waveforms and 400 training samples per cycle. In the example settings listed above, 100 waveforms were tested in each of the 21 cycles; considering a readout speed at video frame rate of 60 Hz, this process corresponds to a training time of ~35 seconds, if not further delayed by the control module and training software.

ACKNOWLEDGEMENTS

Funding: Material preparation, optical characterizations, and device design and fabrication were supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, Materials Sciences and Engineering Division under contract no. DE-AC02-05CH11231 within the Electronic Materials Program (KC1201). Part of the SKM measurements were supported by the U.S. Department of Energy, Microelectronics Science Research Center (Nanoscale Hybrids: A New Paradigm for Energy Efficient Optoelectronics) under contract no. DE. A.O. acknowledges the support of U.S. Army Research Office (ARO). **Author Contributions:** D.Z., A.O., and A.J. conceived the idea for the project and designed the experiments. Y.L., D.Z., and M.M. developed the training algorithms. D.Z., J.G., I.K., and D.U. fabricated the devices. D.Z., J.G., H.M.K., S.W., N.H., and I.K.M.R.R. performed optical measurements. Y.L., D.Z., and S.W. developed the control system. D.Z., T.J.W. and I.K.M.R.R. prepared the MIR filters. D.Z., Y.L., and M.M. analyzed the data. J.G. performed semiconductor device modeling. J.B. helped analyze the MIR photodetection mechanism. D.Z., Y.L., J.B., A.O., and A.J. wrote the manuscript. All authors discussed the results and commented on the manuscript. **Competing financial interests:** The authors A.J and D.Z, affiliated with the University of California, Berkeley, and A.O and Y.L, affiliated with UCLA, are inventors of a United States Provisional Patent Application on the SKM. The remaining authors declare no competing financial interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper or the Supplementary Materials. The data for Figs. 2, 3, 4, and 5 can be found at the data repository Dryad (73).

SUPPLEMENTARY MATERIALS

Supplementary Text
Figs. S1 to S19
Tables S1 to S3
References (54 to 72)
Movie S1




Fig. 1: Architecture of the SKM. (A) The SKM captures a scene with the synchronized operation of a spectral encoder and an electrically tunable bipolar photodetector. The device performs machine learning analysis of the scene's spectra, with the photocurrents directly as the inference results. (B) Reconfigurable computation using a tunable photodetector. Taking a photoconductor as an example, the control voltage $V_{ctr}(t)$ can be applied as the source-drain bias to determine the photo charges' transient travel direction. (C) V_{ctr} controls both the sign and the absolute magnitude of the responsivity $R(\lambda, V_{ctr}(t))$. (D) The data flow during the in-sensor computation. The object's

spectral feature $P(\lambda)$ is first encoded by the spectral encoder $F(\lambda, t)$ at time intervals $t_1, t_2, \dots, t_J$, and then projected to the detector. The encoded light intensity $P(\lambda) \cdot F(\lambda, t)$ is automatically multiplied with $R(\lambda, V_{ctr}(t))$ yielding a transient photocurrent. The current integrates with finite readout bandwidth denoted by the function $\Gamma(t)$. The total photocurrent, $\overline{I_{ph}} = \iint P(\lambda) \cdot F(\lambda, t) \cdot R(\lambda, V_{ctr}(t)) \cdot \Gamma(t) \cdot d\lambda \cdot dt$, can be programmed by $V_{ctr}(t)$ into various kernel machines for task-specific spectral analysis.




Fig. 2. Visible-band SKMs for image segmentation. (A) The ‘sniff’ operation. The SKM is first trained with random pixels in training images with known classes as the ground truth. (B) The trained blue/red $V_{ctr}(t)$ waveforms for selectively highlighting leaf vs. bird regions. Colored shades in the $V_{ctr}(t)$ plot correspond to the synchronized LED colors. (C) The ‘seek’ operation with the trained $V_{ctr}(t)$ waveforms as user commands. (D) The LED spectra for the spectral encoding. (E) The weighted spectrum $Q(\lambda)$ for highlighting the leaf regions. (F) $Q(\lambda)$ for highlighting the bird regions. (G) $\overline{I_{ph}}$ highlighting the leaf regions. (H) $\text{sign}(\overline{I_{ph}})$ of (G), overlying with the raw picture for easier comparison. (I) $\overline{I_{ph}}$ highlighting the bird regions. (J) $\text{sign}(\overline{I_{ph}})$ of (I). (K) SKM combined with back-end spatial post-processing, aiming to harness both the spectral and spatial features of the test sample. A shallow/compact CNN

with a limited number of trainable variables (see fig. S5 for architecture and training details) processes the $\overline{I_{ph}}$ values in (I) to better classify the test image, which resulted in an improved pixel-level classification accuracy of 98.57%.




Fig. 3. SKM results for different tasks. (A) SKM results for wafer oxide thickness identification. The training (upper) and testing (lower) sampling points are with two oxide thicknesses, 90.5 nm (left, red) and 90.0 nm (right, blue). (B) $\overline{I_{ph}}$ values. Each pillar corresponds to one testing point illustrated in (A) with the same color. Regions with 90.5 nm oxide are supposed to show positive current. (C) Photocurrent-based inference of oxide thicknesses. Red diamonds: blind testing experimental readout from testing points on unseen wafers; blue circles: thicknesses and expected readouts of the points from a different batch of training wafers; gray line: expected ($y = x$). (D) Training samples for binary classification of natural and artificial leaves. (E) Blind test results for leaf classification. Colored squares indicate the SKM's classification results, and the letters 'N' and 'A' denote the ground truth labels for natural (positive) and artificial (negative) leaves, respectively. (F) SKM results for leaf hydration level classification. The $\overline{I_{ph}}$ values are plotted against the relative water concentrations of the test leaves never seen before; see the Materials and Methods

section for ground truth measurements. Red dots: low-hydration-level leaves. Blue diamonds: high-hydration-level leaves.




Fig. 4. MIR SKM for chemometrics. (A) Configuration of the gate-tunable bipolar bP-MoS₂ photodiode. (B) The band diagram plotted based on simulations under different gate biases. (C) The responsivity calibrated using a supercontinuum laser directed through a $\lambda = 1.32 \mu\text{m}$ bandpass filter as the light source. (D) The detectivity calibrated with FTIR. (E) Visibly transparent analytes. (F) The total photocurrent for the C₈F₁₈ sample (red line), which was highlighted ($\overline{I_{ph}} > 0$) when the black body IR source is turned ON (white background) and drops back with IR OFF (gray background). The C₂H₆OS sample (blue) was suppressed so that ($\overline{I_{ph}} < 0$) when the IR source is ON. (G) The chart for a complete set of trained binary classifications over four different chemicals in (E). Red: the highlighted materials; blue: the suppressed materials. (H) The total photocurrents, trained at 1 pA every percentage, versus the ground truth water concentration in acetone.




Fig. 5: Modeling of SKM architecture for complex tasks. (A) Crop identification with the Salinas Valley hyperspectral dataset. The visible image is shown on the left. The ground truth of the ‘stubble’ region is highlighted in white over a black background. Further right shows the binary classification results during gradient-blind genetic training on a simulated SKM array. White/black: training pixels identified as positive/negative. Sky blue/navy: test pixels identified as positive/negative. (B) Example RGB image and SKM identification results from a hyperspectral city scene dataset LIB-HSI. (C) Example RGB image and tumor diagnosis results from a hyperspectral human brain dataset. (D) The sensitivity of the classification results in (A)-(C), with a linear-kernel support vector machine (45) run in a digital processor, a transformer (46) run in a digital processor, and an in-sensor SKM. (E) The specificity comparisons. (F) Power consumption for capturing and processing one hypercube of each dataset. (G) The latency for capturing and processing one hypercube.




Fig. 6: Performance evaluations of the SKM architecture. (A) The architecture for hyperspectral analysis in a digital processor, with major data bottleneck at the camera serial interface (CSI). (B) The SKM bypasses the data bottleneck with in-sensor spectral information computation and compression. (C) The speed and resolution in reported works (black dots) and the ideal values for different applications (colored shades). The solid lines correspond to the estimated limits of (A) and (B) set by state-of-the-art imagers and processors (see Supplementary Text for more details). (D) Data throughput and power consumption comparisons.

Corresponding Author

*E-mail: ozcan@ucla.edu, ajavey@berkeley.edu

- ¹ L. Bian, Z. Wang, Y. Zhang, L. Li, Y. Zhang, C. Yang, W. Fang, J. Zhao, C. Zhu, Q. Meng, X. Peng, J. Zhang, A broadband hyperspectral image sensor with high spatio-temporal resolution. *Nature* **635**, 73–81 (2024).
- ⁵ W. Zhang, J. Suo, K. Dong, L. Li, X. Yuan, C. Pei, Q. Dai, Handheld snapshot multi-spectral camera at tens-of-megapixel resolution. *Nat. Commun.* **14**, 5043, (2023).
- ⁶ R. Gebbers, V. I. Adamchuk, Precision agriculture and food security. *Science* **327**, 828-831, (2010).
- ⁷ M. T. Eismann, A. D. Stocker, N. M. Nasrabadi, Automated hyperspectral cueing for civilian search and rescue. *Proc. IEEE* **97**, 1031-1055, (2009).
- ⁸ Y. Jiang, S. Karpf, B. Jalali, Time-stretch LiDAR as a spectrally scanned time-of-flight ranging camera. *Nat. Photonics* **14**, 14-18, (2020).
- ⁹ F. M. Padilla, R. de Souza, M. T. Pena-Fleitas, R. Grasso, M. Gallardo, R. B. Thompson, Influence of time of day on measurement with chlorophyll meters and canopy reflectance sensors of different crop N status. *Precis. Agric.* **20**, 1087-1106, (2019).
- ¹⁰ L. Mennel, J. Symonowicz, S. Wachter, D. K. Polyushkin, A. J. Molina-Mendoza, T. Mueller, Ultrafast machine vision with 2D material neural network image sensors. *Nature* **579**, 62-66, (2020).
- ¹² L. Pi, P. Wang, S. J. Liang, P. Luo, H. Wang, D. Li, Z. Li, P. Chen, X. Zhou, F. Miao, T. Zhai, Broadband convolutional processing using band-alignment-tunable heterostructures. *Nat. Electron.* **5**, 248-254, (2022).
- ¹³ F. Liao, Z. Zhou, B. J. Kim, J. Chen, J. Wang, T. Wan, Y. Zhou, A. T. Hoang, C. Wang, J. Kang, J. H. Ahn, Y. Chai, Bioinspired in-sensor visual adaptation for accurate perception. *Nat. Electron.* **5**, 84-91, (2022).
- ¹⁷ L. Li, S. Li, W. Wang, J. Zhang, Y. Sun, Q. Deng, T. Zheng, J. Lu, W. Gao, M. Yang, H. Wang, Y. Pan, X. Liu, Y. Yang, J. Li, N. Huo, Adaptive machine vision with microsecond-level accurate perception beyond human retina. *Nat. Commun.* **15**, 6261, (2024).

- ¹⁸ Y. Liu, Z. Ji, G. Cen, H. Sun, H. Wang, C. Zhao, Z. L. Wang, W. Mai, Perovskite-based color camera inspired by human visual cells. *Light Sci. Appl.* **12**, 43, (2023).
- ²¹ D. Kwak, D. K. Polyushkin, T. Mueller, In-sensor computing using a MoS₂ photodetector with programmable spectral responsivity. *Nat. Commun.* **14**, 4264, (2023).
- ²² J. Ma, Q. Gu, Y. Liu, J. Lai, P. Yu, X. Zhuo, Z. Liu, J. H. Chen, J. Feng, D. Sun, Nonlinear photoresponse of type-II Weyl semimetals. *Nat. Mater.* **18**, 476-481, (2019).
- ²⁴ D. Zhang, Z. Xu, G. Cheng, Z. Liu, A. R. Gutierrez, W. Zang, T. B. Norris, Z. Zhong, Strongly enhanced THz generation enabled by a graphene hot-carrier fast lane. *Nat. Commun.* **13**, 6404, (2022).
- ²⁵ A. Aizerman, Theoretical foundations of the potential function method in pattern recognition learning. *Autom. Remote Control.* **25**, 821-837, (1964).
- ²⁶ T. Hofmann, B. Schölkopf, A. J. Smola, Kernel methods in machine learning. (2008).
- ²⁷ S. Pai, Z. Sun, T. W. Hughes, T. Park, B. Bartlett, I. A. D. Williamson, M. Minkov, M. Milanizadeh, N. Abebe, F. Morichetti, A. Melloni, S. Fan, O. Solgaard, D. A. B. Miller, Experimentally realized in situ backpropagation for deep learning in photonic neural networks. *Science* **380**, 398-404, (2023).
- ²⁸ A. Momeni, B. Rahmani, B. Scellier, L. G. Wright, P. L. McMahon, C. C. Wanjura, Y. Li, A. Skalli, N. G. Berloff, T. Onodera, I. Oguz, F. Morichetti, P. del Hougne, M. Le Gallo, A. Sebastian, A. Mirhoseini, C. Zhang, D. Markovic, D. Brunner, C. Moser, S. Gigan, F. Marquardt, A. Ozcan, J. Grollier, A. J. Liu, D. Psaltis, A. Alù, R. Fleury, Training of Physical Neural Networks, *ArXiv*, arXiv:2406.03372, (2024)
- ²⁹ A. F. Gad, Pygad: An intuitive genetic algorithm python library. *Multim. Tools Appl.* **83**, 58029-58042, (2024).
- ³⁰ D. Rodrigo, O. Limaj, D. Janner, D. Etezadi, F. J. G. de Abajo, V. Pruneri, H. Altug, Mid-infrared plasmonic biosensing with graphene. *Science* **349**, 165-168, (2015).

- ³² D. Jia, Y. Zhang, Q. Yang, Y. Xue, Y. Tan, Z. Guo, M. Zhang, L. Tian, J. X. Cheng, 3D chemical imaging by fluorescence-detected mid-infrared photothermal Fourier light field microscopy. *Chemical & Biomedical Imaging* **1**, 260-267, (2023).
- ³³ J. Wang, X. Chen, W. Hu, L. Wang, W. Lu, F. Xu, J. Zhao, Y. Shi, R. Ji. Amorphous HgCdTe infrared photoconductive detector with high detectivity above 200 K. *Appl. Phys. Lett.* **99**, (2011).
- ³⁷ M. Amani, E. Regan, J. Bullock, G. H. Ahn, A. Javey, Mid-wave infrared photoconductors based on black phosphorus-arsenic alloys. *ACS Nano* **11**, 11724-11731, (2017).
- ³⁸ C. R. Pidgeon, C. M. Ciesla, B. N. Murdin, Suppression of non-radiative processes in semiconductor mid-infrared emitters and detectors. *Prog. Quantum Electron.* **21**, 361-419, (1997).
- ⁴⁰ N. Higashitarumizu, S. Z. Uddin, D. Weinberg, N. S. Azar, I. K. M. R. Rahman, V. Wang, K. B. Crozier, E. Rabani, A. Javey, Anomalous thickness dependence of photoluminescence quantum yield in black phosphorous. *Nat. Nano* **18**, 507-513., (2023).
- ⁴¹ J. Bullock, M. Amani, J. Cho, Y. Z. Chen, G. H. Ahn, V. Adinolfi, V. R. Shrestha, Y. Gao, K. B. Crozier, Y. L. Chueh, A. Javey, Polarization-resolved black phosphorus/molybdenum disulfide mid-wave infrared photodiodes with high detectivity at room temperature. *Nat. Photonics* **12**, 601-607, (2018).
- ⁴² [https://www.ehu.eus/ccwintco/index.php/Hyperspectral Remote Sensing Scenes](https://www.ehu.eus/ccwintco/index.php/Hyperspectral_Remote_Sensing_Scenes) [Accessed online on Oct 12, 2024]
- ⁴³ N. Habili, E. Kwan, W. Li, C. Webers, J. Oorloff, M. A. Armin, L. Petersson, A hyperspectral and RGB dataset for building façade segmentation. *European Conference on Computer Vision* 258-267, (2022).
- ⁴⁴ H. Fabelo, S. Ortega, A. Szolna, D. Bulters, J. F. Piñeiro, S. Kabwama, A. J. O'Shanahan, H. Bulstrode, S. Bisshopp, B. R. Kiran, D. Ravi, R. Lazcano, D. Madroñal, C. Sosa, C. Espino, M. Marquez, M. D. L. L. Plaza, R. Camacho, D. Carrera, M. Hernández, G. M. Callicó, J. M. Molina, B. Stanciulescu, G. Z. Yang, R. Salvador, E. Juárez, C. Sanz, R. Sarmiento, In-vivo hyperspectral human

brain image database for brain cancer detection. *IEEE Access* **7**, 39098-39116, (2019).

⁴⁵ F. Melgani, L. Bruzzone, Classification of hyperspectral remote sensing images with support vector machines, *IEEE T. Geosci. Remote Sens.* **42**, 1778-1790, (2004).

⁴⁶ D. Hong, Z. Han, J. Yao, L. Gao, B. Zhang, A. Plaza, SpectralFormer: Rethinking hyperspectral image classification with transformers. *IEEE T. Geosci. Remote Sens.* **60**, 1-15, (2021).

⁴⁷ P. Horstrand, R. Guerra, A. Rodríguez, M. Díaz, S. López, J. F. López, A UAV platform based on a hyperspectral sensor for image capturing and on-board processing. *IEEE Access* **7**, 66919-66938, (2019).

⁵¹ Z. Wang, T. Li, R. Du, N. Yang, J. Ping, A high-efficiency lettuce quality detection system based on FPGA. *Comput. Electron. Agric.* **231**, 109978, (2025).

⁵² L. Zhang, C. Zhou, B. Liu, Y. Ding, H. J. Ahn, S. Chang, Y. Duan, M. T. Rahman, T. Xia, X. Chen, Z. Liu, X. Ni, Real-time machine learning-enhanced hyperspectro-polarimetric imaging via an encoding metasurface. *Sci. Adv.* **10**, eadp5192, (2024).

⁵³ Y. Xu, L. Lu, V. Saragadam, K. F. Kelly, A compressive hyperspectral video imaging system using a single-pixel detector. *Nat. Commun.* **15**, 1456, (2024).

⁷³ D. Zhang, Y. Li, J. Geng, H. M. Kim, M. Ma, S. Wang, I. Kim, T. J. Wijaya, N. Higashitarumizu, I. K. M. R. Rahman, D. Urmossy, J. Bullock, A. Ozcan, A. Javey, Spectral kernel machines with electrically tunable photodetectors, Dryad (2025); <https://doi.org/10.5061/dryad.jh9w0vtpv>