A Computer Vision Sensor for AI Accelerated Detection and Tracking of Occluded Objects

Can Cuhadar[1] and Hoi Nok Tsao[1]

[1]Nanyang Technological University

January 11, 2022

Abstract

A prominent problem in computer vision is occlusion, which occurs when an object's key features temporarily disappear behind another crossing body, causing the computer to struggle with image detection. While the human brain is capable of compensating for the invisible parts of the blocked object, computers lack such scene interpretation skills. Cloud computing using convolutional neural networks is typically the method of choice for handling such a scenario. However, for mobile applications where energy consumption and computational costs are critical, cloud computing should be minimized. In this regard, we propose a computer vision sensor capable of efficiently detecting and tracking covered objects without heavy reliance on occlusion handling software. Our edge-computing sensor accomplishes this task by self-learning the object prior to the moment of occlusion and uses this information to "reconstruct" the blocked invisible features. Furthermore, the sensor is capable of tracking a moving object by predicting the path it will most likely take while travelling out of sight behind an obstructing body. Finally, sensor operation is demonstrated by exposing the device to various simulated occlusion events.

Keywords: Computer vision, occlusion handling, edge computing, object tracking, dye sensitized solar cell.

Corresponding author Email: *hoinok.tsao@nie.edu.sg*
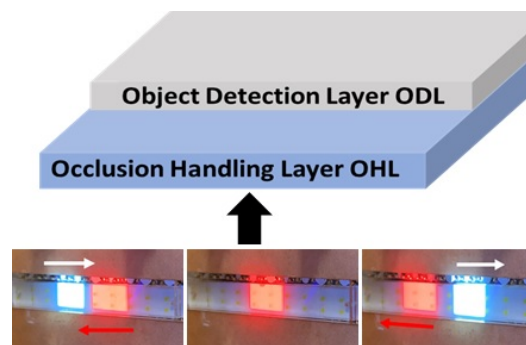
ToC Figure

Figure 1: **ToC Figure.** An AI accelerating computer vision sensor is presented that is capable of detecting and tracking occluded objects.

1

Introduction

The ability to detect objects is vital for autonomous systems such as self-navigating vehicles, security surveillance, or traffic monitoring. While object recognition can be routinely achieved by convolutional neural networks (CNNs), several challenges remain. One of the most severe problem is occlusion, which takes place when features of an object are being masked by other bodies. Such a scenario occurs frequently in traffic, with cars or pedestrians constantly intersecting and passing each other. While the human brain is capable of compensating for the invisible parts of the obscured object, computers lack such scene interpretation skills. To understand why handling occlusion is challenging for machines, we briefly review how current object detection works. State of the art computer vision relies on the following basic concept: An object emits or reflects light that is measured by the corresponding pixels of a camera. The pixels' electric outputs in response to the registered light intensities are commonly known as pixel values. Hence, each object is being perceived by the computer as a distinct stream of pixel values. For object recognition, a computer can be trained to analyse certain key features such as distinct edges or shapes, based on comparing and matching the learned and observed pixel outputs. Such an approach can be very efficiently performed using CNN, currently the most favoured tool(*Gradient-based learning applied to document recognition*, n.d.; *Combining edge and texture information for real-time accurate 3D camera tracking*, n.d.; *ImageNet classification with deep convolutional neural networks — Communications of the ACM*, n.d.; *Edge Detection Using Convolutional Neural Network*, n.d.). However, for CNN to work reliably, a very extensive training dataset is required in order to cover as many different forms that this object can possibly take on(*The Pascal Visual Object Classes Challenge: A Retrospective - International Journal of Computer Vision*, n.d.; "The Pascal Visual Object Classes Challenge: A Retrospective", 2012; "Recent advances in convolutional neural networks", 2018; *ImageNet: A large-scale hierarchical image database*, n.d.). This strong dependence on a complete dataset carries along several limitations. Firstly, the larger the dataset, the more memory is needed, adding to increased costs for running object detection. Secondly, and most importantly, extensive datasets necessitate cloud computing, using the Internet as the vast data source instead of smaller and faster memories localized in the computer itself(*Robust Multi-Modality Multi-Object Tracking*, n.d.; *Revisiting Point Cloud Classification: A New Benchmark Dataset and Classification Model on Real-World Data*, n.d.). As a result, latency in image and video processing may occur, potentially leading to fatal accidents for fast traveling autonomous vehicles that need to continuously monitor traffic in real time for safe self-navigation. These problems are further amplified when dealing with objects that are partially or completely blocked by other objects, causing changes in detected pixel values. As a result, the computer may be tricked into recognizing a different body. Even though this well-known occlusion problem can be mitigated by using an even more extensive training dataset that include images of the object being subjected to all kinds of possible obstructions, such a brute force approach would consume immense energy and memories as well as computational power, all of which are not desirable for always-on and real-time operations such as computer vision(*Occlusion Handling in Generic Object Detection: A Review*, n.d.; "Multiple object tracking: A literature review", 2021).

To solve this problem, sensors are needed that directly process visual data with minimal access to cloud computing. For this purpose, artificial intelligence (AI) accelerating cameras have been developed. These devices contain graphics processing units (GPU) or vision processing units (VPU) that execute image segmentation and classification for CNN to reduce data analysis by the cloud(López et al., 2012). In contrast to these existing AI accelerating cameras, our approach is to integrate video pre-processing directly at the sensor level without the need for energy hungry GPUs or VPUs. Current cameras usually just register light intensities and colour and hence are not able to perform any data analysis by themselves on a hardware level. Our proposed AI accelerating camera should have the following features crucial for efficient detection of "faulty" objects: First, each camera pixel is a single device and does not consist of multiple electronic components or circuitry in order to minimize pixel size for high video resolution and to reduce power consumption. Second, to cut down energy and computational costs, fault tolerant object detection is carried out by the camera itself without the use of any additional computer chips and with minimal runtime of any CNN based software. Current hardware approaches for handling the occlusion problem primarily involve, for instance, stereo vision where multiple cameras operate to cover various angles of the object(*Contour-based*

2

object tracking with occlusion handling in video acquired using mobile cameras, n.d.; *Object Tracking in the Presence of Occlusions via a Camera Network,* n.d.). In this way, there may always exist a perspective in which the occluded body can be detected. Other sensor technologies such as LIDAR or GPS can further render the tracking of occluded objects more robust(*Fusing vision and LIDAR - Synchronization, correction and occlusion reasoning,* n.d.; *Detection scheme for a partially occluded pedestrian based on occluded depth in lidar–radar sensor fusion,* n.d.; *LiDAR Lateral Localisation Despite Challenging Occlusion from Traffic,* n.d.). However, the number of sensors employed increases materials, energy, and computational costs and further creates a heavy load of data to be processed, potentially leading to latency. Therefore, our proposed vision sensor will be a faster and more cost-effective alternative for edge-computed occlusion handling. We will discuss how our sensor deals with various forms of occlusion such as stationary objects crossed by a travelling foreground, or tracking of moving objects blocked by another moving or stationary entity.

Experimental Section/Methods

Device fabrication: Patterned FTO glasses (sheet resistance14 ohm/sq) for sensor fabrication were purchased from Latech. The independence of the TiO_2 photoanodes was achieved by laser etching separate pixel islands on the FTO covered glass. First, patterned FTO glasses were annealed at 500 for cleaning. After cooling down to room temperature, substrates were immersed into a 0.04 M $TiCl_4$ solution and kept in an oven at 70 for 45 min to deposit a compact TiO_2 blocking layer. This step was repeated one more time. Once substrates were cleaned using DI water and EtOH and dried, the 20-pixel and 6 x 4 pixel sensor photoanodes were fabricated by screen printing mesoporous TiO_2 layers of 1mm by 1mm squares using 30NR-D TiO_2 paste (Greatcell Solar), with subsequent annealing at 500 for 30 min. After the substrates were cooled down, they were submerged into a 0.1 mM Dyenamo Yellow (DN-FN01 purchased from Dyenamo and used as is) solution of tert-butanol/acetonitrile (1:1 v:v ratio) over night. Next, the substrates were washed with acetonitrile and sandwiched with another FTO glass as the counter electrode using Surlyn. Finally, the electrolyte was injected between photoanode and counter electrode. The corresponding electrolyte compositions are described in the text. Cobalt complexes were purchased from GreatCell Solar and used as is.

Sensor characterization: Open circuit voltage (V_{OC}) measurements were carried out using a National Instruments NI PXIe-1071 chassis and PXIe-4163 24-channel source measurement unit (SMU). Solis-3C (from Thor Labs) served as the LED white light source for illumination intensity vs time measurements. For the occlusion case studies involving the supplemented movies, an EPSON EH-TW3200 projector was utilized as the light source. The movies were projected onto the sensors by an optical lens. The pixels' V_{OC} outputs were probed simultaneously. For the occlusion of letter 'E', the 24-pixel sensor array was used. To minimize scattered light interfering with the measurement, the sensor was exposed to the letter 'E' by masking the corresponding pixels with black tape. The letter 'E' was projected onto the sensor via a movie until the pixels' V_{OC} reached saturation (for about 5 sec) and subsequently 4 pixels were occluded for 1 sec. After the occlusion, the letter 'E' was projected once more. Same voltage measurement conditions were used for the 24-pixel sensor.

Results and Discussion

1.1. Detection of stationary occluded objects

1.1.1.   *Operation principles of OHL pixels*

Our approach towards managing the occlusion problem is to introduce camera pixels consisting of a two-device tandem structure. The first layer carries out object detection, in the following referred to as object detection layer ODL, utilizing any standard image capturing pixels such as silicon photodiodes. The second layer, which is the focus of this article, handles the correct identification of occluded objects and is hence termed occlusion handling layer OHL, as schematically depicted in Figure 1a. Both sensing layers operate independently, in this way enabling fast parallel computing. We first investigate the simplest scenario involving a stationary object being masked by another body passing by as the foreground (Movie 1 simulates a simple example where a stationary bright square is being blocked by a crossing red square). In this situation, our AI accelerating sensor carries out object detection using the following algorithm (Figure 1b):

Rich media available at https://drive.google.com/file/d/129jvkYHScN1LDZETYDvsEtCCQRkZsyXk/view?usp=sharing
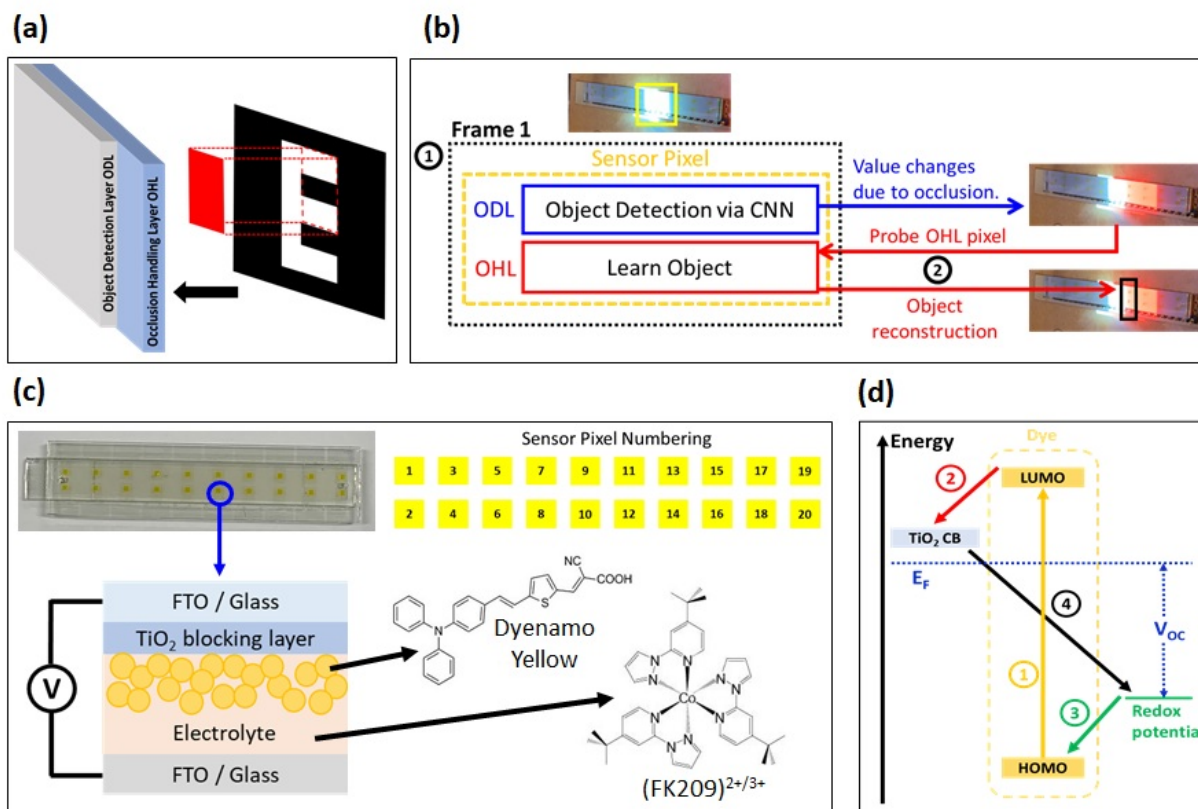
Figure 2: **Figure 1.** This is a caption a) Schematic of a double layer sensor pixel where each layer parallelly and independently handles a different visual task. Here, the object detection layer (ODL) detects the original object whereas the occlusion handling layer (OHL) learns and reconstructs the blocked object. As an example, the parts of the letter 'E" is being masked by a red rectangle. b) Schematic of our sensor's occlusion handling algorithm. Yellow frame in the Frame 1 image indicates detected object by CNN, in this case a white square. Algorithm steps numbered 1 and 2 are detailed in the text. Black frame in image of step 2 highlights the object's reconstructed parts hidden behind a blocking red square foreground. c) Image, pixel numbering and device architecture of our proof-of-concept OHL consisting of 20 pixels. FTO stands for fluorine doped tin oxide. The yellow spheres represent TiO_2 nanoparticles sensitized with photoactive dyes. d) Charge transfer mechanisms occurring in the OHL pixel during illumination.

Step 1: The object of interest is first identified by running CNN on the ODL pixels' electric outputs in the first image frame. For all successive image frames, the ODL pixels are being continuously measured without triggering CNN. As long as these pixel values do not change, the same object is being perceived. This approach significantly reduces CNN runtime, especially at high frame rates for stationary or slow-moving objects. At the same time, the OHL pixels register and memorize the light intensities corresponding to the detected object (self-learning step).

Step 2: Once occlusion occurs, the ODL pixel values change. At this instance, the OHL pixel outputs are probed and should give the same values prior to the occlusion event, in this way reconstructing the visually blocked original object without having to run CNN. Hence, our sensor allows simple distinction between the occluded background (performed by the OHL) and the blocking foreground (detected by the ODL) without having to run computationally expensive cloud computing or other occlusion handling software.

To achieve such image reconstruction under occlusion, the OHL pixels have to maintain their electric outputs,

and hence "recall" the image, when the light intensity or colour change due to the appearance of a different body masking the original object. We will implement such unsupervised self-learning as an invariance in the pixels' open circuit potential V_{OC} once light conditions are altered. This feature can be attained by dye-sensitized solar cells (DSSC) via the appropriate electrochemical modifications. Another advantage of using a DSSC is the possibility of good transparency, thus allowing the desired two functional layer tandem architecture. The DSSC based OHL pixels consist of a mesoporous transparent TiO_2 layer containing chemically anchored photoactive dyes. An electrolyte infiltrates this photo sensitive film. The counter electrode allows connection of the cell to an external load for signal extraction (Figure 1c). Various charge transfer mechanisms take place during light exposure. In the following, only those processes majorly involved in the V_{OC} build-up and decay will be summarized (Figure 1d)(Grätzel, 2003; *Photoelectrochemical cells - Nature*, n.d.). Process 1: The photoactive dye molecule absorbs photons and excites an electron from the highest occupied molecular orbital (HOMO) to the lowest unoccupied molecular orbital (LUMO). Process 2: The excited electron within the LUMO is injected into the TiO_2 photoanode conduction band CB (charge injection). The dye is now missing an electron; it is in its oxidized state. In order to continue absorbing photons and injecting electrons into the TiO_2 photoanode, this oxidized dye molecule regains an electron back from a nearby reducing agent in the electrolyte. This process 3 is called dye regeneration. The reducing agent now becomes an oxidizing species after having transferred an electron to the oxidized dye. Some injected electrons residing within the TiO_2 photoanode may transfer to these nearby oxidized shuttles to transform them back to reducing agents. This electron loss, process 4, is called charge recombination and will decrease the charge carrier density within the photoanode, lowering the TiO_2 layer's quasi Fermi level E_F. As a consequence, the V_{OC} decreases, since it is defined by the potential difference between E_F and the redox potential of the redox agents(Grätzel, 2003; *Photoelectrochemical cells - Nature*, n.d.; *Dye-Sensitized Solar Cells: Fundamentals and Current Status - Nanoscale Research Letters*, n.d.; "Effects of driving forces for recombination and regeneration on the photovoltaic performance of dye-sensitized solar cells using cobalt polypyridine redox couples", 2011). Normally, in the dark, all electrons accumulated and stored in the TiO_2 film during illumination will rapidly discharge through recombination, immediately resetting the V_{OC} close to zero, thus losing any memory effect. In the following, we use the transition from bright to dark to simulate the most extreme change in light intensity that could occur during occlusion. A more typical real life occlusion event would rather involve more moderate alterations in light intensities when two different objects intersect.
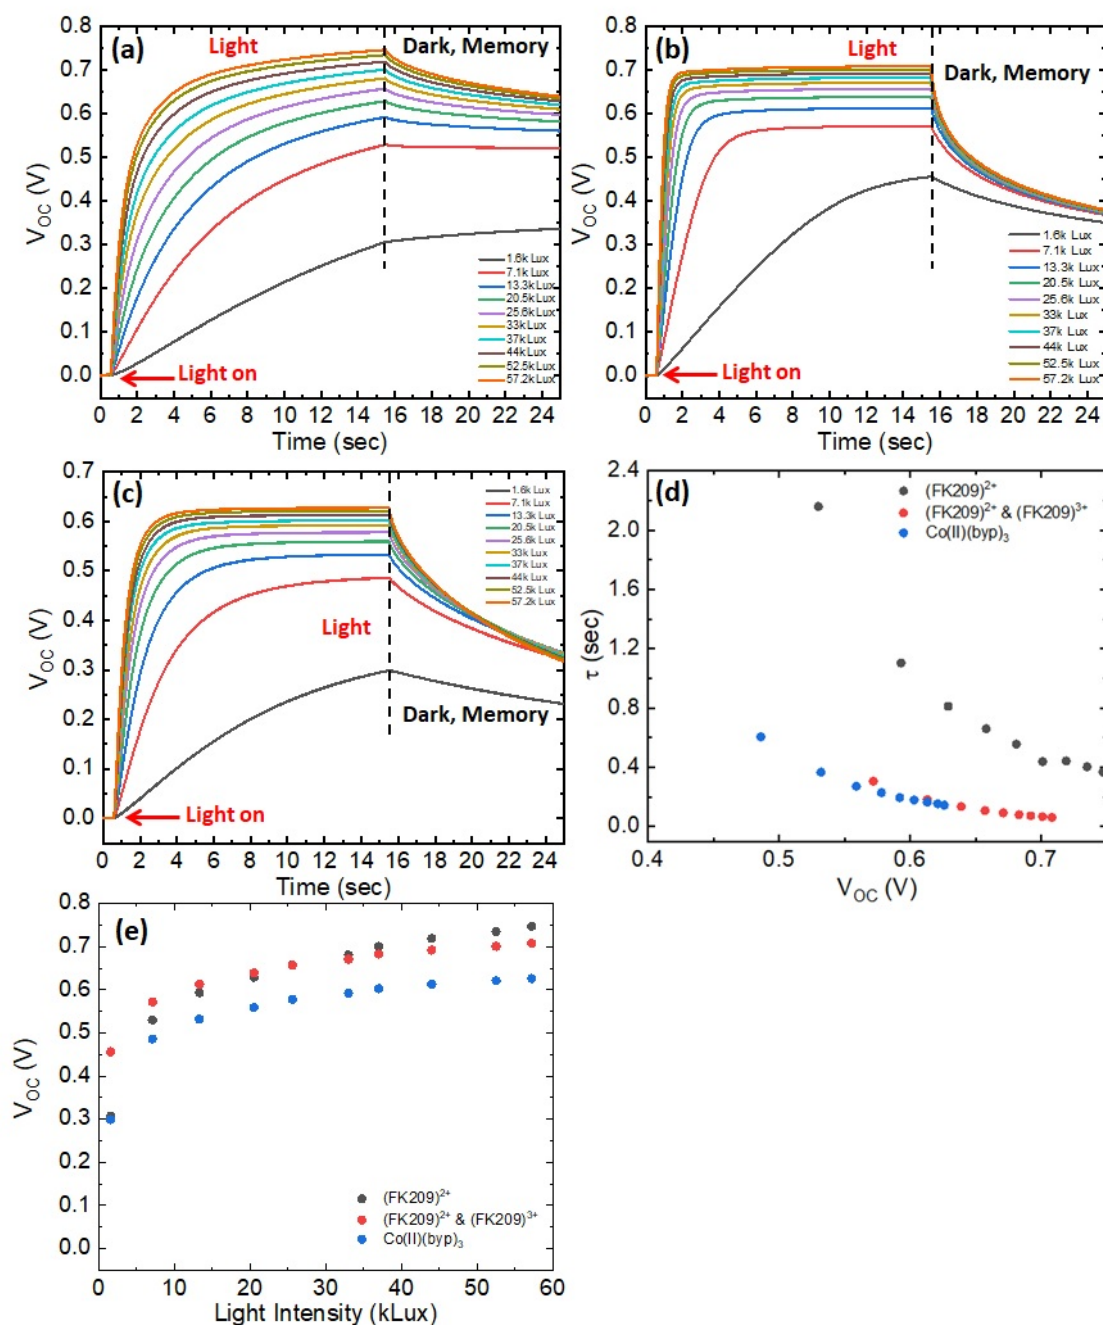
Figure 3: **Figure 2.** a) V_{OC} response with illumination time at various LED white light intensities of a single occlusion handling pixel, containing $(FK209)^{2+}$ only, b) both $(FK209)^{2+}$ and $(FK209)^{3+}$, and c) $Co(II)(bpy)_3$ only redox couples. Black dashed vertical line indicates the moment when switching from light to dark (from object detection to memory). Hence, the portion of the graphs right of this dashed line represents the pixel's learning behaviour. d) TiO_2 electron lifetimes at various V_{OC} for all three pixels corresponding to graphs a), b), and c). e) V_{OC} vs light intensity for a single OHL pixel containing various redox couple compositions.

7

The OHL pixels consist of yellow organic dye (Dyenamo yellow) sensitized mesoporous TiO_2 photoanodes infiltrated by a liquid methoxypropionitrile electrolyte containing 0.1 M Tris(2-(1H-pyrazol-1-yl)-4-tert-butylpyridine)cobalt(II)di(bis(trifluoromethylsulfonyl)imide), or in short $(FK209)^{2+}$ (*Regeneration and recombination kinetics in cobalt polypyridine based dye-sensitized solar cells, explained using Marcus theory*, n.d.) as the reducing agent (Figure 1c). No oxidizing species were added in order to minimize the presence of recombination centers to which the TiO_2 electrons would readily transfer, resulting in high recombination rates and hence potential loss of memory feature due to a rapid V_{OC} decay(*Regeneration and recombination kinetics in cobalt polypyridine based dye-sensitized solar cells, explained using Marcus theory*, n.d.; Ondersma, 2013; *Dye-Sensitized Solar Cells: Driving-Force Effects on Electron Recombination Dynamics with Cobalt-Based Shuttles*, n.d.). To probe the learning behaviour of this sensor, the V_{OC} response of a single OHL pixel was measured when switching from light to dark at various intensities (Figure 2a). It is evident that the V_{OC} does not rapidly drop to zero but decays at a slow rate, showing that learning can be achieved. More precisely, for all probed intensities, the V_{OC} retains more than 90% of the initial value for a period of 10 s in the dark. To further investigate the importance of recombination rate on the learning behaviour, DSSC pixels were fabricated containing both $(FK209)^{2+}$ as the reducing agent and its corresponding oxidizing species Tris(2-(1H-pyrazol-1-yl)-4-tert-butylpyridine)cobalt(III)tris(bis(trifluoromethylsulfonyl)imide), in short $(FK209)^{3+}$, with the latter acting as additional recombination centers to enhance recombination. As Figure 2b reveals, this device's V_{OC} declines much faster upon switching from light to dark, clearly indicating a much shorter memory capacity. That is, this pixel loses more than 10% of its initial V_{OC} value after a significantly shorter period of time (around 1 s) for all probed intensities. The rate of V_{OC} exponential decay in the absence of light is a method for quantifying the recombination rate since, at open circuit in the dark, all the electrons residing in the TiO_2 photoanode have to transfer back to the oxidizing species in the electrolyte to reach an energetically equilibrated state(*The Influence of Charge Transport and Recombination on the Performance of Dye-Sensitized Solar Cells*, n.d.). Using this open circuit voltage decay method, the TiO_2 electron lifetimes τ (the reciprocal of charge recombination rate) with respect to V_{OC} were obtained for both types of pixels (Figure 2d and Table 1)("Determination of rate constants for charge transfer and the distribution of semiconductor and electrolyte electronic energy levels in dye-sensitized solar cells by open-circuit photovoltage decay method", 2004; "Dye-sensitized solar cells: Driving-force effects on electron recombination dynamics with cobalt-based shuttles", 2010; *Dye-Sensitized Solar Cells: Driving-Force Effects on Electron Recombination Dynamics with Cobalt-Based Shuttles*, n.d.). Evidently, the TiO_2 electron lifetimes in the $(FK209)^{2+}$ only pixel are much longer than in the device containing the oxidizing species $(FK209)^{3+}$, proving that the period of memory retention is critically linked to the TiO_2 charge recombination rate. However, one has to keep in mind that charge recombination can be dependent on the dye regeneration rate as well. In especially, when a reducing agent regenerates the dye less efficiently, fewer oxidizing shuttles will be formed that can pose as charge recombination centers. To verify this point, DSSC pixels containing 0.1 M Tris-(2,2'-bipyridine)cobalt(II)di(bis(trifluoromethane)sulfonimide), in short $Co(II)(bpy)_3$, were fabricated since this cobalt complex yields faster regeneration of the Dyenamo Yellow dye than $(FK209)^{2+}$ ("Effects of driving forces for recombination and regeneration on the photovoltaic performance of dye-sensitized solar cells using cobalt polypyridine redox couples", 2011). As illustrated in Figure 2d, the TiO_2 electron lifetime is indeed shorter than the $(FK209)^{2+}$ pixels, triggering a much faster V_{OC} drop when transitioning from light to dark (Figure 2c as compared to Figure 2a). In fact, the $Co(II)(bpy)_3$ based pixel exhibits similar V_{OC} and TiO_2 electron lifetimes as the $(FK209)^{2+}$ & $(FK209)^{3+}$ device (Figure 2d and Table 1), despite the lack of initial oxidizing species $Co(III)(bpy)_3$ acting as additional recombination centers in the former cell. This observation implies that, due to the faster dye regeneration rate of $Co(II)(bpy)_3$, the amount of the oxidizing shuttle $Co(III)(bpy)_3$ being generated during illumination is comparable to the initial concentration of oxidizing species in the $(FK209)^{2+}$ & $(FK209)^{3+}$ device, ultimately leading to the observed similar recombination rates. It is to be noted that, for all three types of pixels discussed this far, the V_{OC} decay in the dark is most retarded after exposure to low light intensities (1.6k Lux to 7.1k Lux in Figure 2a, b, and c). The reason lies in the reduced number of charge carriers injected into and residing in the TiO_2 photoanode at low intensities, leading to diminished V_{OC} and a decreased driving force for recombination("Recombination and

8

transport processes in dye-sensitized solar cells investigated under working conditions", 2006; "A numerical model for charge transport and recombination in dye-sensitized solar cells", 2006). Furthermore, fewer dyes are active at dimmer conditions, resulting in more suppressed dye regeneration which in turn curtails the number of oxidizing species being created. Consequently, less recombination centers are present, yielding prolonged TiO_2 electron lifetimes (slower recombination rates) at lower V_{OC} (Figure 2d).

Intensity (Lx)	$(FK209)^{2+}$		$(FK209)^{2+}$ & $(FK209)^{3+}$		Co^{2+} (bpy)	
	V_{OC} (V)	τ (sec)	V_{OC} (V)	τ (sec)	V_{OC} (V)	τ (sec)
1.6k	0.31	-4.21	0.46	1.49	0.3	4.21
7.1k	0.53	2.16	0.57	0.31	0.49	0.61
13.3k	0.59	1.1	0.61	0.18	0.53	0.37
20.5k	0.63	0.81	0.64	0.13	0.56	0.27
25.6k	0.66	0.66	0.66	0.11	0.58	0.23
33k	0.68	0.56	0.67	0.09	0.59	0.2
37k	0.7	0.44	0.68	0.08	0.6	0.18
44k	0.72	0.44	0.69	0.07	0.61	0.16
52.5k	0.74	0.4	0.7	0.07	0.62	0.15
57.2k	0.75	0.37	0.71	0.06	0.63	0.14

Table 1: **Table 1.** TiO_2 electron lifetimes τ of single pixels cells containing $(FK209)^{2+}$, $(FK209)^{2+}$ & $(FK209)^{3+}$ and $Co(II)(bpy)_3$ electrolytes corresponding to Figures 2a, b, c in the main text, respectively.

1.1.2.   Case studies of occluded stationary objects

To assess how well the OHL can manage occlusion of a stationary object, a proof-of-concept 2 x 10 pixel sensor was fabricated (image in Figure 1c), containing the $(FK209)^{2+}$ only electrolyte. The sensor is exposed to the aforementioned situation illustrated in movie 1, involving a simple white 2 x 2 pixel square temporarily blocked by a crossing 2 x 2 pixel red square (Figure 3a), simulating a typical occlusion scenario. Ideally, for the sensor to correctly manage occlusion, the OHL pixels' V_{OC} responses corresponding to the stationary white square should not deviate much while the blocking red square passes by. Figure 3b shows these four active pixels' temporal responses while detecting this scene. The white square's pixels' maximum V_{OC} drop is less than 10% during the whole occlusion event (indicated by the blue and green dashed frames in Figure 3b), implying that the sensor is indeed capable of reconstructing the masked object. Note that without the pixels' learning effect, this V_{OC} decay would immediately decay close to zero during the intersection by the dimmer red square foreground, as evidenced by the much lower V_{OC} in response to red light (red line in Figure 3b). thus failing to recognize the covered white square.

To further demonstrate the occlusion handling capabilities of our AI accelerating camera, we expanded the case study to the detection of a more complex object, using a proof-of-concept sensor consisting of 4 x 6 OHL pixels. The letter 'E' was chosen as the stationary image being briefly occluded for 1s by a black 2 x 3 pixel rectangle such that 'E' temporarily appears as the letter 'L' (Figure 3c). The V_{OC} outputs of the OHL pixels seeing the letter 'E' were recorded simultaneously and are depicted in Figure 3c. The four blue curves reflect the responses of those four pixels that are directly exposed to the blocking black rectangle. During the occlusion event, marked by the two dashed vertical red lines, these four pixels experience the largest deviation in V_{OC} (the inset illustrates the absolute percent changes in V_{OC} for each sensor pixel, whereas

Figure 3f summarizes the corresponding absolute V_{OC} values, with the four occluded pixels outlined by the blue frames). However, these deviations are small as compared to any standard ODL pixel when transitioning from light to dark. Therefore, our sensor is able to correctly recognize the original object 'E' by reconstructing the blocked key features that may have tricked the computer to see 'L' instead. More precisely, the ODL detects the letter 'L', while the OHL allows the computer to interpret the situation to be actually the original letter 'E' being obstructed by a black rectangle. To handle a new appearing object, the OHL pixels' memory has to be erased via a reset to zero V_{OC}. This can be achieved by simply short circuiting the cells, in this way depleting all the stored electrons in the TiO_2 film.

So far, we have considered only cases where the occlusion occurs by a darker foreground. Generally, feature masking can also take place by a brighter object. The OHL pixels' responses to a sudden increase in light intensities are summarized in Figures 3d and 3e, highlighting that the pixels exhibit V_{OC} retention of more than 76% for the probed region of light exposure. Within the limit of these light intensity ranges, occlusion can be handled through a variation of pixel value of less than 24%.
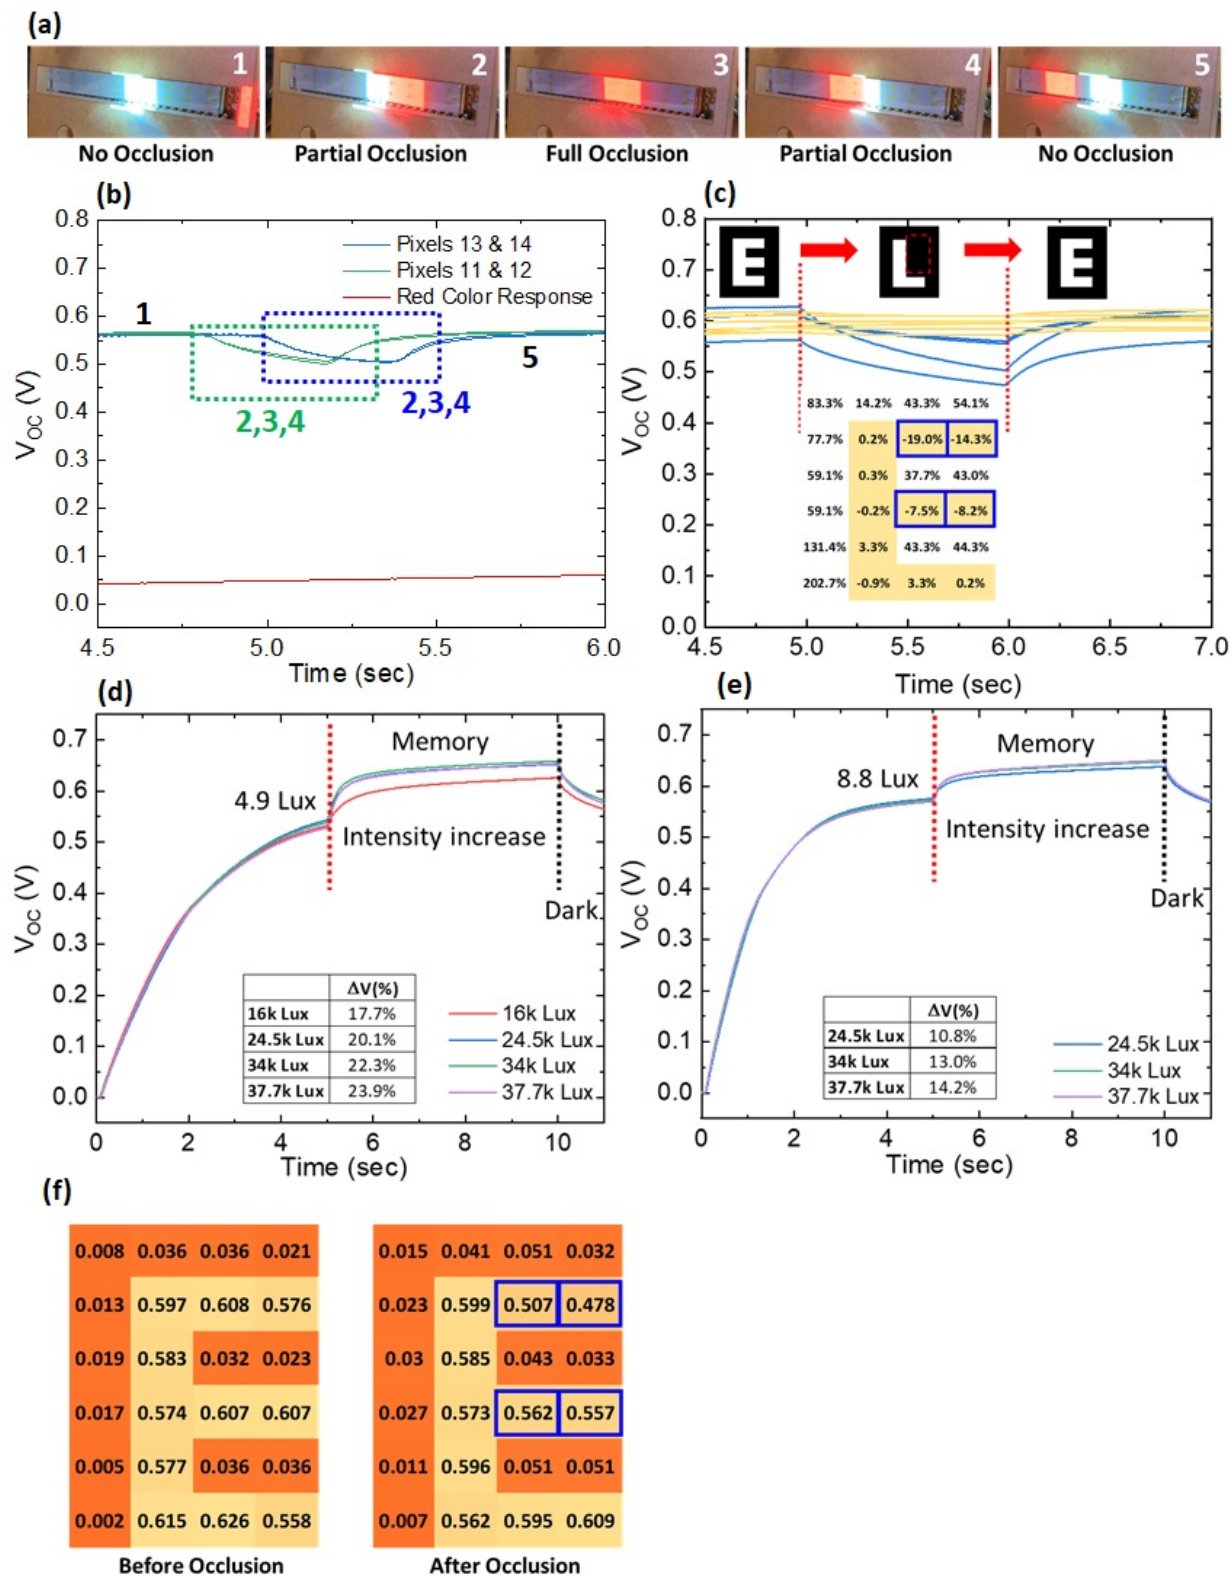
(a)

1	2	3	4	5
No Occlusion	Partial Occlusion	Full Occlusion	Partial Occlusion	No Occlusion

(b)

Pixels 13 & 14
Pixels 11 & 12
Red Color Response

V_{OC} (V) vs Time (sec)

1   2,3,4   2,3,4   5

(c)

E → L → E

83.3%	14.2%	43.3%	54.1%
77.7%	0.2%	-19.0%	-14.3%
59.1%	0.3%	37.7%	43.0%
59.1%	-0.2%	-7.5%	-8.2%
131.4%	3.3%	43.3%	44.3%
202.7%	-0.9%	3.3%	0.2%

(d)

Memory, 4.9 Lux, Intensity increase, Dark

	ΔV(%)
16k Lux	17.7%
24.5k Lux	20.1%
34k Lux	22.3%
37.7k Lux	23.9%

16k Lux, 24.5k Lux, 34k Lux, 37.7k Lux

(e)

Memory, 8.8 Lux, Intensity increase, Dark

	ΔV(%)
24.5k Lux	10.8%
34k Lux	13.0%
37.7k Lux	14.2%

24.5k Lux, 34k Lux, 37.7k Lux

(f)

Before Occlusion:

0.008	0.036	0.036	0.021
0.013	0.597	0.608	0.576
0.019	0.583	0.032	0.023
0.017	0.574	0.607	0.607
0.005	0.577	0.036	0.036
0.002	0.615	0.626	0.558

After Occlusion:

0.015	0.041	0.051	0.032
0.023	0.599	0.507	0.478
0.03	0.585	0.043	0.033
0.027	0.573	0.562	0.557
0.011	0.596	0.051	0.051
0.007	0.562	0.595	0.609

Figure 4: **Figure 3.** a) Snapshots of the stimulated occlusion experiment where a white stationary square is obstructed by a moving red square. b) V_{OC} responses with time of the four OHL pixels 11, 12, 13, 14 (pixel position according to Figure 1c) of our proof-of-concept 20-pixel OHL sensor perceiving a white 2 x 2 pixel square being temporarily blocked by a red square of the same size moving from right to left, as depicted by the image frames 1-5 in a). The corresponding moments of occlusion numbered 2, 3, and 4 in a) are indicated by the green and blue dashed frames as seen by pixel pairs 11, 12 and 13, 14, respectively. The red curve depicts the OHL pixel's V_{OC} response to red light exposure time. c) V_{OC} response with time of our proof-of-concept 6 x 4 pixel OHL sensor detecting the letter E being temporarily occluded by a black rectangle (indicated by the dashed red frame) to the letter L. The inset depicts the percent changes in the OHL pixels'

1.2. Tracking of moving occluded objects

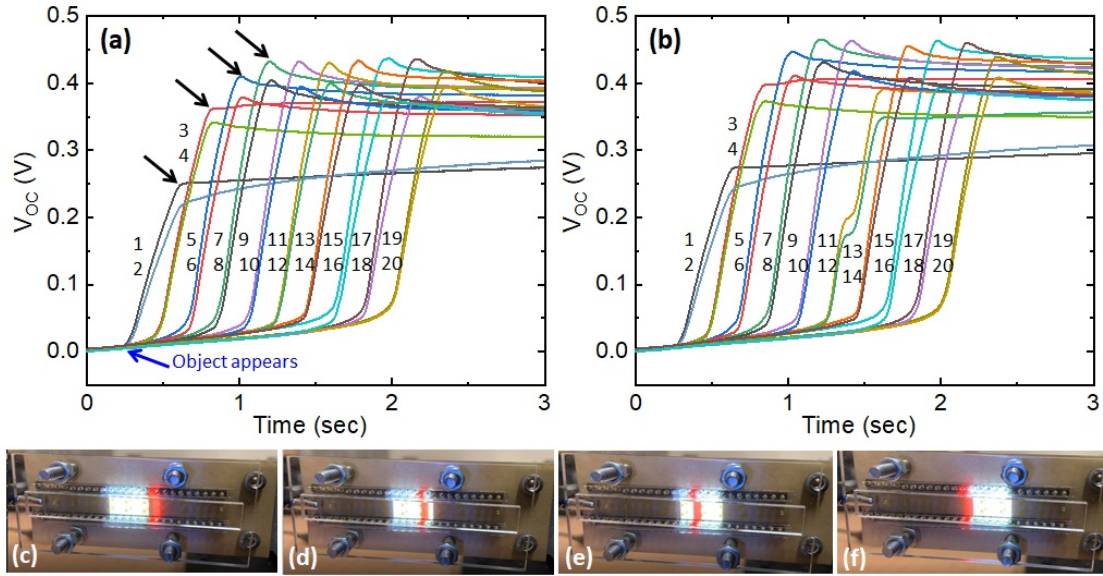### 1.2.1. Sensor requirements for tracking occluded objects

Figure 5: **Figure 4.** a) Simultaneous V_{OC} responses of all 20 OHL pixels exposed to a 2 x 2 pixel white square travelling from left to right at a uniform speed of 50 ms per image frame, encountering no occlusion. Pixel numbering follows Figure 1c. The black arrows (only a few shown as examples) indicate the instances in which the object has passed the corresponding OHL pixel, leading to a change in registered light intensity (from bright to dark in this case). This occurs at the corresponding pixel's peak V_{OC}. b) V_{OC} response of all 20 OHL pixels exposed to a 2 x 2 pixel white square travelling from left to right and occluded by a crossing 2 x 1 red vertical strip, both moving at the same uniform speed as in a). c) Image frame of a square moving from left to right, right before occlusion by a thin red strip takes place. d) e) Image frames during occlusion by a thin red strip. f) Image frame after the occlusion event.

Rich media available at https://drive.google.com/file/d/1LSHrxRdjjI1cn4FvzS1yNSQ2oBonPGV-/view?usp=sharing

Self-navigating systems such as cars or drones necessitate tracking of objects for traffic monitoring or surveillance. The thus far presented operation principle of our occlusion handling vision sensor for stationary objects does not always equally hold for a moving body. To further illustrate this point, let's first consider a simple scenario where the intersecting foreground is small in size as compared to the tracked object. In especially, we again investigate a 2 x 2 pixel white square traveling from left to right at a uniform speed and occluded by a 2 x 1 pixel thin red vertical strip moving in the opposite direction, using our 20-OHL pixel proof-of-concept sensor (movie 2). Figures 4a and 4b depict the sensor's response during the entire white square's motion without and with the thin strip foreground, respectively. Most pixels values in Figure 4b show only minor deviations from the occlusion-free reference scenario reflected in Figure 4a, implying that the sensor

successfully tracked the object despite occlusion. If this was not the case, the pixels seeing the portion of the white square being blocked by the red strip should deliver a significant drop in V_{OC} pertaining to the pixels' red light responses (Figure 3b). The absence of such dramatic V_{OC} decay upon occlusion is due to the pixels' memory effect. This image reconstruction via self-learning by the sensor was only possible because, at any instant of time, the OHL pixels were able to perceive the travelling object prior to the crossing red strip foreground, as highlighted by the snapshots in Figures 4d and 4e.

Rich media available at https://drive.google.com/file/d/1d2KlLFMixozqdctUOLOWuFZjpQO3W7iO/view?usp=sharing

On the other hand, let's consider the case where some OHL pixels are occupied by a stationary obstructing foreground and hence cannot detect the moving object at any time. An everyday analogy would be a camera tracking a car that passes by and disappears behind another vehicle parked in the foreground. Those pixels seeing the foreground never had the opportunity to learn the car image and hence will not be able to recognize it. This situation is simulated by a 2 x 2 pixel white square travelling from left to right while occluded along the way by a stationary 2 x 2 pixel red square, as illustrated in movie 3, with Figures 5a-e mirroring snapshots of this event. For the image frames in which occlusion takes place (Figures 5b to d), those OHL pixels representing the masked portion of the object have never seen and hence were not able to learn these hidden parts and are hence not able to reconstruct them. To manage such a scenario, path prediction is a viable method which can be performed by utilizing the OHL pixels' learning function. That is, tracking can be achieved by forecasting the route the obstructed object will most likely take based on its previous trajectory. Our human brain works in a similar fashion, anticipating an occluded object to reappear following its past motion, based on experience. It is a natural assumption we subconsciously make. To foresee the most likely trail, the OHL pixels should be able to detect and to learn the direction and speed of the path taken by the object.

1.2.2. *How OHL pixels estimate speed of a moving object*

To attain speed estimation, the OHL pixels are engineered in such a way that their V_{OC} output is dependent on illumination time. This approach is based on the concept that light from fast objects dwell on the corresponding pixels for a shorter time than slower ones. As Figure 2a illustrates, the pixel's V_{OC} slowly rises with illumination time. This analog V_{OC} response provides temporal information about the detected motion, i.e. at a given light intensity, the V_{OC} pixel values reveal the speed of the moving object. To demonstrate this type of object tracking, for simplicity, we have limited ourselves to one dimensional motion first. Expansion to two-dimensional motion can be achieved by integrating more pixels. Figure 5f depicts the V_{OC} values of each pixel registering a 2 x 2 pixel white square moving from left to right at a uniform frame rate of 50 ms (movie 4). Since the motion is uniform in our experiment, the V_{OC} responses of each pixel should be similar at a given speed. Furthermore, various constant speeds should result in different V_{OC} corresponding to the analog V_{OC} dependence on light exposure time. This expected behavior is in alignment with our results, as illustrated in Figure 5f, with faster motion yielding lower V_{OC} values (statistical interpretation can be seen on Figure 5i). However, it is to be noted that the OHL pixels themselves cannot correctly perform speed estimation since one V_{OC} value may represent different speeds (illumination times) depending on the registered light intensity (Figure 2a). Hence, to properly assess the speed, both V_{OC} (from OHL) and light intensity values (from ODL) have to be known. This requirement is satisfied by the double layered sensor architecture (Figure 1a). Therefore, our sensor, in addition to its self-learning capabilities, also allows very straightforward speed estimation by simply probing both the ODL and OHL pixels without having to measure time between successive image frames representing the object's motion, leading to significantly reduced computational costs.

Rich media available at https://drive.google.com/file/d/1t8bCnbdeVeWvaJ88WWms4sHgPYFlZoxU/view?usp=sharing
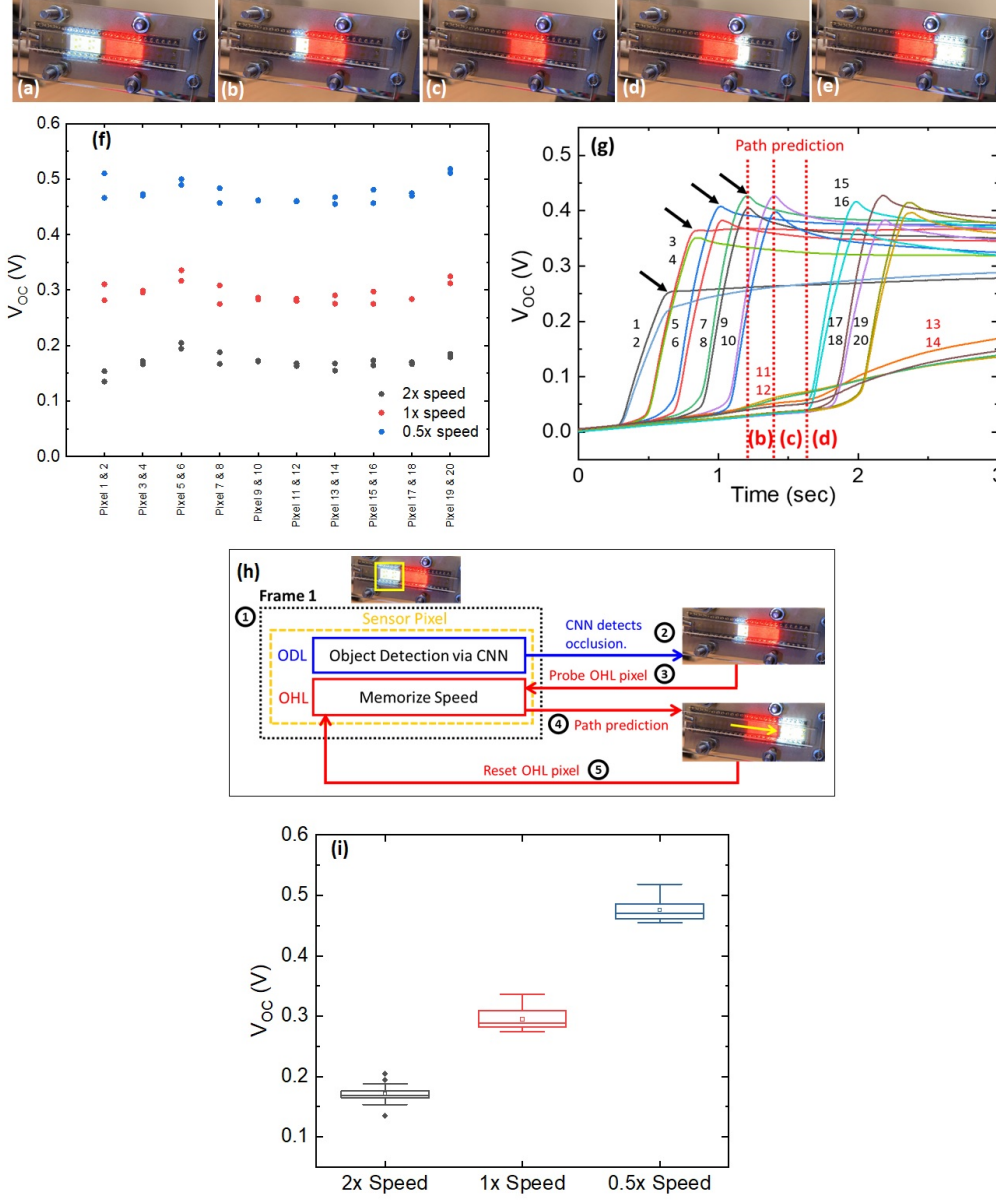
Figure 6: **Figure 5.** a) to e) Representative image frames of a white 2 x 2 pixel square moving from left to right at 50 ms per frame and being b) partially to c) completely occluded by a stationary 2 x 2 pixel red square. f) V_{OC} response of our proof-of-concept 2 x 10 pixel sensor tracking a 2 x 2 pixel white square moving at various uniform speeds. g) V_{OC} response of our proof-of-concept sensor tracking the moving white square undergoing occlusion as illustrated in Figures a) to e). Pixel numbering for the corresponding curves follow Figure 1c. The time regions marked by the dashed vertical red lines refer to the according scenarios b), c), and d). The pixels 11-14, marked red, are exposed to the stationary red foreground. Black arrows (just a few presented as examples) indicate the V_{OC} values at the instance the object has past and is not being perceived anymore by the corresponding pixels. h) Occlusion handling algorithm, with the numbered instruction steps described in detail in the text. i) Box plot corresponding to Figure 5f.

14

The analog V_{OC} response with illumination time is attributed to a gradual accumulation of electrons in the TiO_2 film caused by an increase in dye regeneration rate with light exposure duration. This hypothesis is supported by the general trend that the V_{OC} rises more slowly at lower light intensities for all three types of pixels previously discussed (Figures 2a, b, and c). Let's consider the devices containing no initial oxidizing agents (Figure 2a and 2c). At low intensities where dye excitation is minimal, only a few oxidizing species are being generated via dye regeneration. These Co^{3+} oxidizing shuttles are too low in concentration to diffuse from the mesoporous TiO_2 film into the bulk electrolyte("Mass Transport of Polypyridyl Cobalt Complexes in Dye-Sensitized Solar Cells with Mesoporous TiO2 Photoanodes", 2008). As a consequence, the Co^{3+} compounds are localized near the dye molecules, prohibiting fast dye regeneration by the surrounding Co^{2+} reducing complexes(Tsao, 2018). As illumination continues, more Co^{3+} manage to diffuse from the dye molecules, facilitating the majority Co^{2+} complexes to participate in dye regeneration. This process increases the Co^{3+} concentration with time, leading to gradually enhanced Co^{3+} diffusion from the mesoporous TiO_2 film into the bulk Co^{2+} electrolyte, and thus increasing the dye regeneration rate. Consequently, the number of electrons injected into the TiO_2 rises with illumination time, resulting in the observed analog V_{OC} growth.

1.2.3. *Case studies of tracking moving occluded objects*

Now that we have established how speed estimation works, we come back to our initial occlusion scenario, simulating a simple one-dimensional motion of a white square travelling from left to right and encountering blockage by a stationary red square (Figure 5a-e and movie 3). Figure 5g shows all 20 pixels' V_{OC} values during the whole event (pixel numbering and positions follow Figure 1c). Clearly, the pixels seeing the stationary red square (pixels 11-14, marked red) deliver low V_{OC} throughout the whole occlusion scenario (Figures 5b, c, d); hence, failing to reconstruct parts of the travelling white square disappearing behind the red one. As mentioned before, one way to solve this problem is for the sensor to predict the object's path during occlusion based on the past trajectory taken. This can be achieved by simply probing the self-learning OHL pixels 1-8 that have tracked the object prior to the onset of occlusion (Figure 5b). As Figure 5g reveals, pixels 1-8 maintain mostly similar V_{OC} values after changes in light intensity (indicated by black arrows) throughout the whole event (Table 2), hence providing a speed and path prediction that allows the sensor to "assume" the object keeps following the original trajectory behind the obstructing red square. More precisely, during the period in which half of the white square vanishes (marked as (b) in Figure 5g and corresponding to Figure 5b), probing the pixels' 1-8 would provide information about the object's previous speed and past route. Similarly, for complete occlusion (Figure 5c), pixels 9 and 10 also have to be measured within the period of occurrence, designated by (c) in Figure 5g. The camera can then keep tracking the temporary disappearing object along this predicted path. It is open for the user to decide on the length of the past course that should be considered. In our example, we chose the whole track represented by pixels 1-10. However, for some scenarios, a shorter trail involving fewer pixels may be sufficient as well.

Pixel	1	2	3	4	5	6	7	8	9	10
ΔV_{OC} (%)	0.01	0.04	0.01	-0.02	-0.02	-0.0	-0.03	-0.04	-0.03	-0.03

Table 2: **Table 2.** Percent changes in V_{OC} (ΔV_{OC} (%)) of pixels 1-10 corresponding to Figure 5g, probed during partial (marked as section (b)) and complete occlusion (marked as section (c)). Changes were calculated using the differences between the V_{OC} at the instance of light intensity alteration and the V_{OC} at the end of the corresponding occlusion scenario (b) and (c).

The following algorithm summarizes in more detail how the sensor tracks moving occluded objects (Figure 5h):

Step 1: The object of interest is first identified by running CNN on the ODL pixels' electric outputs in the

first image frame. For all successive image frames, the ODL pixel outputs are being continuously measured without triggering CNN. As long as these pixel values do not change, the same object is being perceived. This approach significantly reduces CNN runtime, especially at high frame rates for stationary or slow-moving objects.

Step 2: When ODL pixel values change, CNN has to determine which of the following scenarios is taking place: a) A new object has entered the scene, b) the object of the previous frame started to move, or c) the object of the previous frame is occluded. In the following, we focus on the case of occlusion as the other possibilities a) and b) can be handled at low cost by a locally trained CNN without having to invoke cloud computing.

Step 3: If occlusion occurs in the new frame, the computer probes the OHL pixels previously occupied by the travelling object (in our simple scenario, pixels 1-8 will be measured).

Step 4: These OHL pixel values together with the object's corresponding ODL pixel outputs provide speed estimation. The positions of these active pixels reveal the spatial information required for complete path prediction.

Step 5: The probed OHL pixels are being reset by short circuiting the devices, allowing analysis of a new object. The user would need to define the condition for pixel resetting. One reasonable requirement could be resetting after the tracked object is confirmed by CNN to have reappeared after the occlusion event (Figure 5e).

Rich media available at https://drive.google.com/file/d/1eeI6U1-6t0UQFWp_5ZqAdb7FLHIBSPoB/view?usp=sharing

The last typical occlusion scenario we would like to explore is when both the background and the foreground are moving and crossing each other, a situation that occurs frequently among pedestrians during crowd monitoring for instance. To illustrate how our device copes with such a case, we again exposed our proof-of-concept 20-pixel sensor to the same 2 x 2 pixel white square travelling from left to right and encountering a 2 x 2 pixel red square moving in the opposite direction, acting as the occluding entity. Figure 6a to 6e are snapshots of this scenario, with both objects moving at the same speed of 50 ms per frame (movie 5). Figure 6f depicts the corresponding sensor pixels' responses during this whole event. The occlusion handling algorithm presented in Figure 5h applies for this situation as well. At the onset of partial overlap (Figure 6b), the sensor triggers the computer to probe pixels 1-8 in order to perform path prediction (steps 2 and 3 in Figure 5h). During complete occlusion (Figure 6c), pixels 9 and 10 have to be measured in addition. As Figure 6f illustrates, all pixels show only minor deviations in V_{OC} upon a sudden change to darkness (Table 3), revealing that correct tracking of the occluded object is indeed possible via recollection of its previous trajectory.
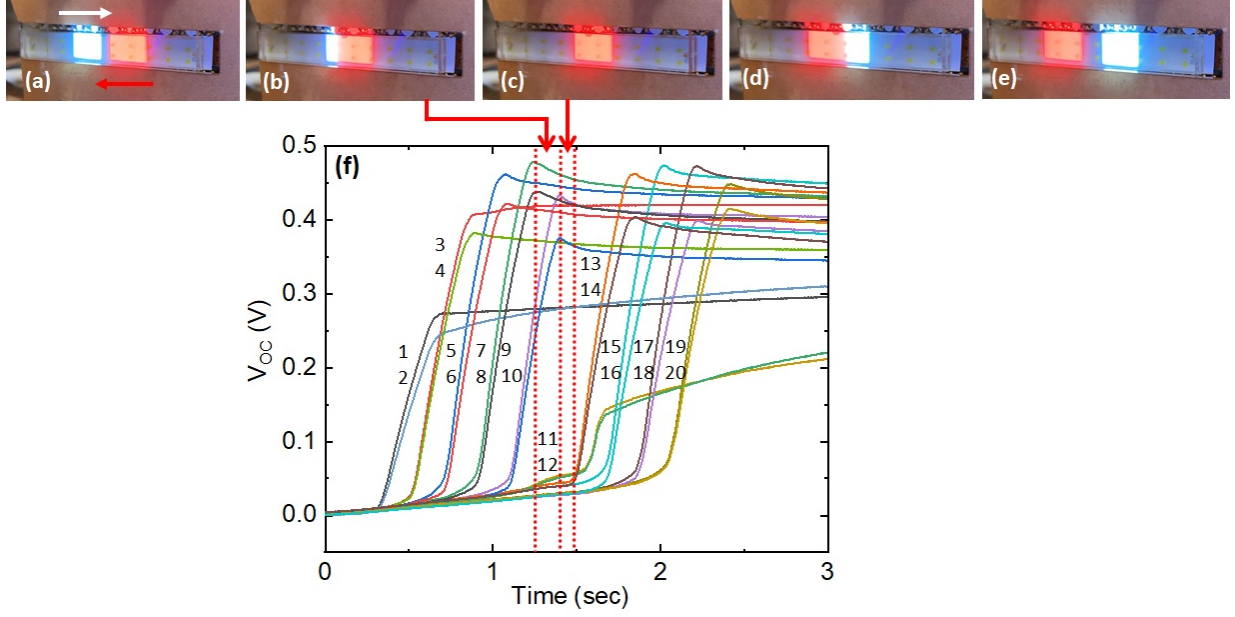
Figure 7: **Figure 6.** a) to e) Representative image frames of a white 2 x 2 pixel square moving from left to right at 50 ms per frame and being b) partially to c) completely occluded by a crossing 2 x 2 pixel red rectangle. f) V_{OC} response of our proof-of-concept 20-pixel sensor tracking this moving occluded square. Red vertical dashed lines indicates the time periods of partial b) and complete occlusion c).

Pixel	1	2	3	4	5	6	7	8	9	10
ΔV_{OC} (%)	0.01	0.04	0.01	-0.01	-0.02	-0.01	-0.02	-0.01	-0.01	-0.01

Table 3: **Table 3.** Percent changes in V_{OC} (ΔV_{OC} (%)) of pixels 1-10 corresponding to Figure 6f and probed during onset of occlusion (Figure 6b). Changes were calculated using the differences between the V_{OC} at the instance of light intensity alteration and the V_{OC} at the end of the corresponding occlusion scenario marked by the dashed red vertical lines.

Conclusion

In conclusion, we have demonstrated a video pre-processing sensor that acts as a convolutional neural network accelerator by handling occlusion directly at the intelligent system's level, leading to cost effective and latency-free real-time object detection and tracking. The core idea towards achieving such an edge-computing computer vision camera is to introduce photovoltaic pixels that exhibit a dual function. First, the pixels are capable of learning the perceived object and utilizes this knowledge to reconstruct it when being blocked, without the use of any computationally expensive and potentially laggy cloud computing. Second, the pixels have the ability to estimate the speed of a moving object via an analog open circuit voltage dependence on illumination time. This feature enables tracking of an obstructed object by predicting its most probable path based on the route travelled prior to the occlusion event. We have illustrated how our sensor successfully handles some simple occlusion scenarios related to everyday situations. Expanding the number of pixels would potentially allow the detection and tracking of more complex objects travelling in different directions.

17

Acknowledgements

Conflict of interest

The authors declare no conflict of interests.

References

https://ieeexplore.ieee.org/document/726791. https://ieeexplore.ieee.org/document/726791

https://ieeexplore.ieee.org/document/1383042. https://ieeexplore.ieee.org/document/1383042

https://dl.acm.org/doi/10.1145/3065386. https://dl.acm.org/doi/abs/10.1145/3065386

https://link.springer.com/chapter/10.1007/978-3-319-40663-3_2. https://link.springer.com/chapter/10.1007/978-3-319-40663-3_2

https://link.springer.com/article/10.1007%2Fs11263-014-0733-5. https://link.springer.com/article/10.1007/s11263-014-0733-5

(2012). *International Journal of Computer Vision*, *111*.

(2018). *Pattern Recognition*, *77*.

https://ieeexplore.ieee.org/document/5206848. https://ieeexplore.ieee.org/document/5206848

https://ieeexplore.ieee.org/document/9010755. https://ieeexplore.ieee.org/document/9010755

https://arxiv.org/abs/1908.04616. https://arxiv.org/abs/1908.04616v2

https://ieeexplore.ieee.org/abstract/document/9378657. https://ieeexplore.ieee.org/abstract/document/9378657

(2021). *Artificial Intelligence*, *293*.

Comparison of GPU-based methods for handling point cloud occlusion. (2012). *Spanish Computer Graphics Conference (CEIG)*, *September*.

https://ieeexplore.ieee.org/document/1335457. https://ieeexplore.ieee.org/document/1335457

https://ieeexplore.ieee.org/document/4379711. https://ieeexplore.ieee.org/document/4379711

https://ieeexplore.ieee.org/document/5548079. https://ieeexplore.ieee.org/document/5548079

https://www.spiedigitallibrary.org/journals/optical-engineering/volume-56/issue-11/113112/Detection-scheme-for-a-partially-occluded-pedestrian-based-on-occluded/10.1117/1.OE.56.11.113112.full?SSO=1. https://www.spiedigitallibrary.org/journals/optical-

engineering/volume-56/issue-11/113112/Detection-scheme-for-a-partially-occluded-pedestrian-based-on-occluded/10.1117/1.OE.56.11.113112.full

https://ieeexplore.ieee.org/document/9109911. https://ieeexplore.ieee.org/document/9109911

Dye-sensitized solar cells. (2003). *Journal of Photochemistry and Photobiology C: Photochemistry Reviews*, *4*.

https://www.nature.com/articles/35104607. https://www.nature.com/articles/35104607

https://nanoscalereslett.springeropen.com/articles/10.1186/s11671-018-2760-6. https://nanoscalereslett.springeropen.com/articles/10.1186/s11671-018-2760-6

(2011). *Journal of Physical Chemistry C*, *115*.

https://pubs.rsc.org/en/content/articlelanding/2013/CP/c3cp50997d. https://pubs.rsc.org/en/content/articlelanding/2013/cp/c3cp50997d

Recombination and redox couples in dye-sensitized solar cells. (2013). *Coordination Chemistry Reviews*, *257*.

https://pubs.acs.org/doi/abs/10.1021/la904643t. https://pubs.acs.org/doi/abs/10.1021/la904643t

https://chemistry-europe.onlinelibrary.wiley.com/doi/10.1002/cphc.200800708. https://chemistry-europe.onlinelibrary.wiley.com/doi/10.1002/cphc.200800708

(2004). *Journal of the American Chemical Society*, *126*.

(2010). *Langmuir*, *26*.

(2006). *Journal of Physical Chemistry B*, *110*.

(2006). *Journal of Physical Chemistry B*, *110*.

(2008). *The Journal of Physical Chemistry C*, *112*.

Illumination Time Dependent Learning in Dye Sensitized Solar Cells. (2018). *ACS Applied Materials & Interfaces*, *10*.