Marshalling Signals Dataset of Radar and Event-Based Camera for Sensory Fusion

Leon Müller[1,2], Manolis Sifalakis[1], Sherif Eissa[2], Amirreza Yousefzadeh[1],
Paul Detterer[1], Sander Stuijk[2], and Federico Corradi[2]
[1]Stichting IMEC Nederland, High Tech Campus 31, Eindhoven 5656 AE, The Netherlands
[2] Eindhoven University of Technology, Department of Electrical Engineering
Email: manolis.sifalakis@imec.nl, f.corradi@tue.nl

Abstract—**The advent of neural networks capable of learning salient features from variance in the radar data has expanded the breadth of radar applications, often as an alternative sensor or a complementary modality to camera vision. Gesture recognition for command control is probably the most commonly explored application. Nevertheless, there is a lack of suitable benchmarking datasets to assess and compare the merits of the different proposed solutions. Furthermore, most current publicly available radar datasets used in gesture recognition lack diversity or generality and are not challenging enough. We make available a unique dataset designed to meet these objectives, combining two synchronized modalities: radar and dynamic vision camera for experimenting with sensory fusion. Moreover, we propose a sparse encoding of the time domain (ADC) signals that achieve a dramatic data rate reduction (>76%) while retaining the efficacy of the downstream FFT processing (<2% accuracy loss on recognition tasks). Finally, we demonstrate early sensory fusion results based on range-Doppler maps from this radar data encoding versus the conventional approach, and dynamic vision, achieving higher accuracy than either modality alone.**

Index Terms—**Radar Full Body Gestures Recognition, FMCW Radar, Raw Radar Data, Sensory Fusion**

I. Introduction

Modern radar systems achieve high-resolution and high precision in range, speed, and angle-of-arrival estimation thanks to GigaHertz frequency bands, high-performance Analog to Digital Converters (ADC), and a large number of antennas. Next, using neural networks to learn salient features from variance in the radar data has expanded the breadth of applications where radar is used, often as an alternative sensor or a complementary modality to camera vision. Radar sensing has several advantages compared to vision; it is more privacy-preserving and effective even in occluded or light-limited conditions. Gesture recognition for command control is probably the most commonly explored [1]–[7], albeit not the only one (e.g. air handwriting [8], agent tracking [9], drone altitude control [10], human gait recognition [11], object detection [12]).

Motivated by explorations with low-power AI-based processing of radar data for edge applications (empowered by DNN and SNN accelerators [13]–[17]), and the intent to combine radar data with other modalities such as the DVS camera [18] in edge AI applications, we created a unique dataset, that can serve the needs of benchmarking purposes, and sensor-fusion applications. It combines dynamic vision with low-resolution single-antenna radar sensing [19] so as to make the task non-trivial; it uses complex enough whole-body gestures; and has enough diversity of data samples to challenge the AI application model generalizability. Arguably, for many low-power edge-AI applications, high data precision is often not needed. It comes at the cost of a deluge of data volumes, which increases the computational, memory, and communication cost and can make radar technology challenging to apply in resource-constrained environments. Most applications use the compressed-sensing Fourier transformed data (range maps, range-Doppler maps, μDoppler maps) to confine the amount of data fed from the sensor to downstream processing. This requires buffering of several chirps in memory, followed by floating-point processing. Here, based on the collected dataset, we looked at how to alleviate this problem of on-radar memory storage and processing, and we propose a binary encoding of the raw radar data. It is cheap to implement in SW or HW in the sensor and preserves all the frequency spectrum information required in commonplace downstream processing. The remaining of this paper is organized as follows. In Methods (section II), we describe our hardware setup of the two modalities for generating the marshaling signals dataset and introduce the encodings and networks we used in our experiments. In Results (section III), we demonstrate the efficiency of the event-based radar signal encoding and its use with the dataset for sensor fusion. In Discussion IV we provide further insights and contextualize our work in the radar and AI literature.

II. Methods

A. Hardware Platform

We developed a data capturing platform that consisted of two sensing modalities: vision with a DVS camera from IniVation [20], and radar based the 8GHz UWB FMCW SISO radar sensor from Imec [19] (Table I). The two modalities have been synchronized at *ms* resolution using hardware signaling by feeding the rising edge of the PRI pulse signal of the radar as an event in the event stream of the DVS. This enabled us to timestamp and enumerate the radar chirps using the clock of the DVS camera. In this way, we are able in real-time to associate and align individual radar chirp transmissions with the corresponding set of polarity events generated by the DVS camera. The underlying idea is illustrated in Fig.1a.

(a) SISO 8GHz radar (top) and DVS camera (bottom).

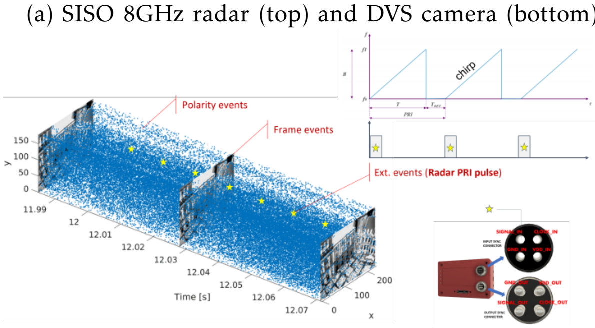

(b) Hardware-based data synchronization of the sensors.

Figure 1: Data capturing platform

The platform is controlled by a software tool developed in the open-source C++ OpenFrameworks[1], for configuring the two sensors, switching on their synchronization, controlling data capturing, and visualizing the data in real-time.

B. Multisensory Aircraft Marshalling Signals Dataset

Using the multisensor platform discussed above, we recorded a dataset that consists of 10 whole-body gestures performed by multiple subjects in different indoor locations (hence different sets of reflections) at 8 discrete distances from the sensor platform ($1m$-$4.5m$, in steps of $0.5m$). The training set features recordings by 6 different subjects in 2 different locations. The testing set contains an additional 7 subjects not present in the training set and a third location. The 10 gestures come from the aircraft marshalling signal code: *turn right*, *turn left*, *straight ahead*, *stop engines*, *start engines*, *slow down*, *move back v1*, *move back v2*, *move ahead*, *emergency stop*, and we have added a background *no-signal* class with people walking or empty space (no gesture being performed).

Parameter	Value
Center Frequency (GHz)	7.3
Bandwidth (MHz)	750
Chirp Time (us)	40.96
Time between Chirps (ms)	1.3
Chirps per frame	192
ADC Samples per Chirp	512
Sampling Frequency (MHz)	12.5
Range Max. (m)	76.288
Range Res. (m)	0.149
Velocity Max. (m/s)	15
Velocity Res. (m/s)	0.5

Table I: FMCW 8 GHz SISO Radar specifications

This dataset is unique (compared to other published datasets) in the following aspects. Firstly, it combines two synchronized event-based sensor modalities: DVS and radar. It is therefore suitable for experimentation with sensor fusion models. Secondly, it is rich in spatio-temporal content for event-based processing: each gesture occupies a long time scale (order of seconds) and is repeated at several distances from the sensors. Thirdly, it has sufficient diversification, which is essential for ML model generalization: each gesture is repeated by several subjects in 2 spaces of different geometry and lighting conditions, and at two different angles from the orthogonal plane to the sensor beam.

Other essential characteristics of this dataset are that the selected gestures from the marshaling signal code have sufficient complexity to engage more than one of the fundamental features of the radar sensing (so a single feature, e.g., distance or radial velocity, or 1 direction of movement, are not sufficient to classify a gesture). And, the fact that the radar employed is low-resolution (8Ghz), with a single-antenna, makes the task of recognition across distances and in face of room reflections more challenging. Finally, we make available the raw ADC data from the radar to maximize its usefulness in experimentation.

Table II, provides a 5-point comparison of this dataset with a few other popular datasets in the literature of experimentation with gestures.[2]

[1] https://openframeworks.cc/

[2] A pre-published version of the dataset for review purposes of this manuscript is available upon request

Dataset	Modalities	Type	# Gestures	# Environments	# People
IBM DVS128 [21]	DVS	full body	11	3	29
ASL-DVS [22]	DVS	hand	24	1	5
IMEC Gestures [7]	8 GHz Radar	hand	4	1	1
Google Soli [23]	60 GHz FMCW Radar	hand	11	1	10
DopNet [24]	24 GHz FMCW Radar	hand	4	1	6
Ours	8 GHz FMCW Radar + DVS	full body	10 + (none)	3	13

Table II: Overview of related datasets for radar gestures and our dataset

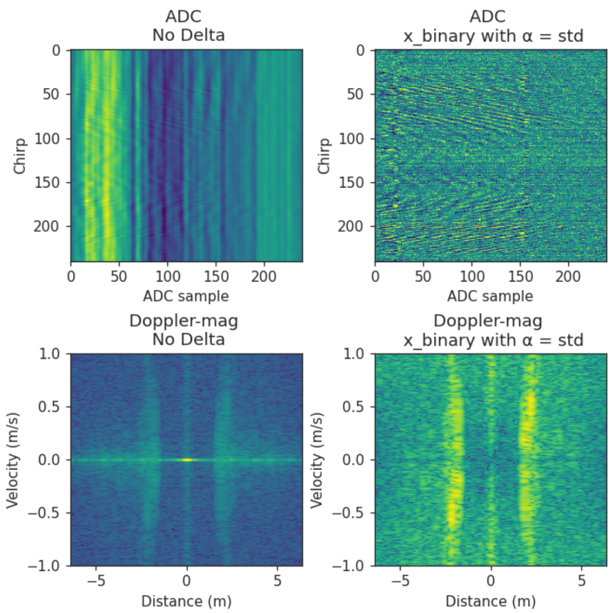

Figure 2: Comparing the impact of our encoding on the ADC data (top) and Doppler-mag maps (bottom). On the left, we see the raw ADC data and the resulting Doppler-mag map. On the right, we show binary ADC data (x_{binary}) where $\alpha = \sigma$ and the resulting Doppler-mag map.

C. Baseline neural networks models

In our experiments, we used the publicly available neural network models EfficientNet-B1 and ResNet18 for the classification of the radar data alone, the DVS data alone, and the fusion of the radar and DVS data. We have used pre-trained models on the ImageNet classification challenge, and we have fine-tuned them to perform classification tasks on our dataset. For the fusion of the two modalities we used one of the three (RGB) input channels of the model for the radar data and the other two for the DVS data (positive versus negative polarity events).

D. Radar and DVS signal representations

For the experiments reported in section III we preprocess the data of the two modalities as follows.
For the radar modality the first encoding we use is the usual transformation of the raw ADC data with the 2D FFT transform to generate range-Doppler maps. We generate frames of chirps (= 960 radar chirps) that correspond to 1.2s duration (enough to include a complete gesture), and then we apply a Han window followed by FFT chirp-wise and followed by another FFT across the chirp dimension. Finally, we crop the map to only the included range and Doppler bins up to a distance slightly longer than the max distance in the dataset. (see Fig. 2)
The second radar encoding is similar to the first only this before applying the FFT transforms, we first apply a first-order delta filter on the ADC data between chirps (slow time), then do level-crossing (thresholding), and everything below the threshold is set to zero, everything above the threshold is set to +1/-1 depending on the sign of the delta. Mathematically this is equivalent to

$$x_{delta}(n,m) = x(n,m) - x(n-1,m)$$
$$x_{binary}(n,m) = \begin{cases} 1, & \text{if } x_{delta}(n,m) > \alpha \\ -1, & \text{if } x_{delta}(n,m) < -\alpha \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

where n is the n-th chirp, m is the m-th ADC sample, and we set the value of α to the standard deviation across a single capture, which is largely the same in the same environment but varies in different environments. The delta operation high-pass filters the data removing static objects, while the thresholding operation removes noise (and numerical rounding errors). We then generate range-Doppler maps from this very sparse representation of the ADC data.
Finally, for the DVS modality, we accumulate the events corresponding to the timespan of each radar frame (= 960 chirps) into a single DVS frame, which creates a flow representation of the gesture movement. Then we normalized the frame in the range $[-1, 1]$. (see examples in the top row of Fig. 3).

III. Results

A. Efficiency and informativeness of binary FMCW radar representation

The new binary encoding of the ADC data proposed in this work is a much sparser set of signals

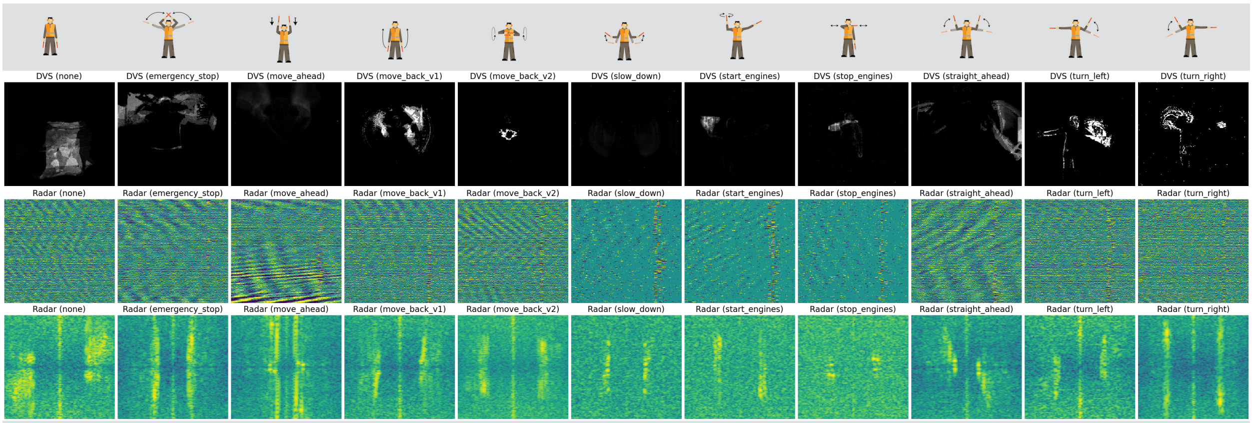

Figure 3: Examples of the gestures of the dataset: DVS (top), x_{binary} processed ADC data (middle) and resulting Doppler-Mag frames (bottom).

than the direct time-domain ADC output. The combination of the delta operation with level-crossing makes the signal drastically more sparse. The binarization operation creates an extreme saving in memory for storing the chirp data: we only need 1-bit to represent a binarized data point, compared to 8/16/32-bit storage for integer values (in the case of raw ADC data). We report an average reduction of about 76% in our experiments' data volumes. Moreover, we can also reduce the computational cost of the FFT transforms in this way, through zero-skipping, and eliminating multiplications. On the other hand, the range-Doppler maps that we obtain from these bitmaps, surprisingly at first impression, preserve all the frequency information as the "conventional" FFT DSP, sufficient to complete the classification task (see Fig. 2). There is an amount of ambient noise introduced in the range-Doppler maps, which is directly related to, and controlled by, the threshold level in level-crossing.

B. Aircraft Marshaling Signals Classification

We demonstrate performance in a classification task using DVS data alone versus radar data alone, and the sensor fusion of the two modalities by combining DVS and radar data. We train the models by leaving one or more people out, and then we used them in the test set (leave-one-out validation). For the DVS-only case, the fine-tuned the (initialized with ImageNet weights) EfficientNet-B1 and ResNet18 can reach an all-class accuracy of 82.7% and 74.6% respectively, after just five training epochs. Similarly, for the radar-only case, we finetuned ResNet18 and EfficientNet-B1 using image inputs representing the Doppler-Mag as visible in Fig. 2. A comparison of the impact of binarized ADC data encoding is reported in Fig. 5. The left set of accuracies corresponds to range-Doppler maps generated from the raw ADC data ($e = none$), and the middle set corresponds to range-Doppler maps from the binarized ADC data with the threshold set to zero ($e = binary$). The right set of accuracies corresponds to range-Doppler maps generated from the binarized ADC data and is further denoised with threshold $\alpha = \sigma$ ($e = lvl{-}x, binary$). Overall, accuracy is minimally affected (even positively with a non-zero threshold), with most results being within 2% of one another. At the same time, the ADC data's density reduction is 41.39% when the threshold is set to zero and 72.76% when the threshold is set to $\alpha = \sigma$. With an ADC density of only 27.24%, the accuracy with all gestures has increased by 1.27%, going from 63.31% to 64.58%. And when considering only a subset of the best-5 performing classes, the accuracy is increased by 2.07% from 94.08% to 96.15%. When we fuse together the two modalities, we can clearly see a boost in the overall accuracy (table III) outside of the statistical variation of either modality separately, which confirms that the information given to the model from each modality is partly orthogonal to the other. This is also confirmed in Fig. 3, by looking at the 3rd (*move ahead*) and 6th (*slow down*) gestures, where the DVS frames have very little event content, but the radar frames of the binarized ADC data show a distinctive texture pattern for each gesture. Finally, as we see in Fig. 4 there are some of the gestures that are similar enough in their information content in both modalities, such as to be confused by the model, thereby confirming the complexity inherent in the dataset (in this case increasing the time span of the data frames helps disambiguate the classes).

IV. Discussion

Because of the chirp-level synchronization of the two modalities in our capturing setup, the herein

Networks			Test Accuracy (All Gestures)			Test Accuracy (Best 5)		
Name	**Parameters**	**Architecture**	**Radar**	**DVS**	**Fusion**	**Radar**	**DVS**	**Fusion**
ResNet18	11.2M	18 layers	59.1	74.6	82.4	86.9	91.7	95.8
EfficientNet-B1	6.5M	11 blocks	64.6	82.6	86.7	96.1	94.8	98.7

Table III: Overview of the classification accuracy using DNN models by sensory modality.

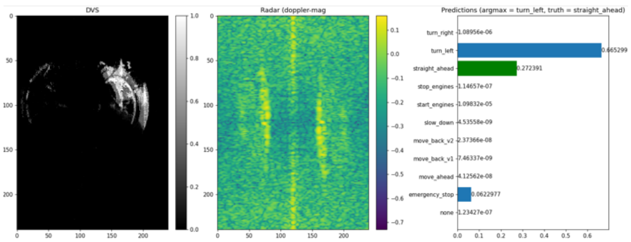

Figure 4: Example where straight_ahead gestures have been wrongly classified as turn_left by EfficientNet-B1. DVS frame (left), radar range-Doppler from binary ADC data x_{binary} (middle), and classifier predictions (right).

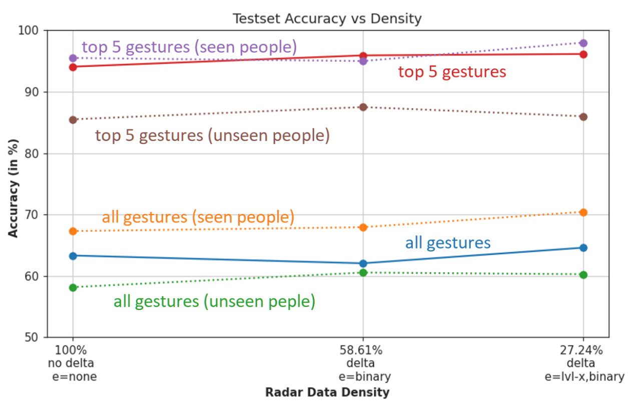

Figure 5: Classification accuracy with EfficientNet-B1 on the three radar encoding methods. The first column uses the Doppler-Mag on full raw ADC data. The second column uses delta encoding with $\alpha = 0$ ($e = binary$). The third column uses level-crossing and binary output with $\alpha = \sigma$ ($e = lvl - x, binary$). The density is greatly reduced in the latter two cases (42.4%, 72.8%), while the classification accuracy is minimally affected.

published dataset offers the possibility to derive data points for training ML models, that have finer resolution than a *radar frame* (which is almost exclusively used in the literature). This opens new perspectives for experimentation and exploration, such as sequence models and recurrent architectures in general that can deal with lower dimensional data in a streaming fashion. Moreover having access in particular to *multimodal temporal* raw data from radar and DVS is very appealing as it enables to explore the temporal dimensions of recurrent neural network models, for both traditional deep learning and spiking neural networks [1]–[5], [25].

The herein presented binary encoding of the ADC data also offers the possibility for exploring and developing more closely integrated sensor-neural network solutions that target resource efficiency and also likely energy efficiency. In a somewhat similar spirit, the authors in [2] proposed the use of the Moving Target Indicator (MTI) operator before the FFT pipeline (typically MTI follows the FFT processing). This is very similar to the delta operator with level crossing which we use in our binary encoding but is more complex (it needs to compute a running average for smoothing, which is, anyway, a linear operation that the neural network can easily incorporate). This similarity provides insight into the results we showed in section III, namely that applying FFT after delta/level-crossing/binarization has *insignificant* information loss. Leaving aside the thresholding and, in our case, additionally, the binarization operations, since MTI (or delta) and FFT are both linear operations, they are communicative, and their order can thus be interchanged without changing the combined effect on the data. Next to the binarization operation, another critical difference between our encoding and the preprocessing in [2] is that we apply the delta operation between subsequent chirps, not frames, which as more minimalistic affects the memory footprint and temporal resolution preserved in the event encoding.
Finally, our binary ADC data encoding for radar due to its simplicity, can be cheaply implemented at the radar digital back-end right after the ADC, or even more efficiently at the analog front-end drastically reducing the quantization resolution of the ADC.

V. Conclusion

We have introduced a new dataset with synchronized prime forms of radar (raw ADC) and DVS camera (polarity events) data based on whole-body gestures from the aircraft marshaling signals. Additionally, we have proposed a lean low-overhead sparse binary encoding for the radar signal returns that dramatically reduce the data rate for FMCW radars without sacrificing their frequency content. This method can be easily implemented after the ADC sampling stage (in dedicated HW logic or in SW). Finally, we demonstrated both the use of the dataset and the encoding with a sensor-fusion classification task using commonplace fine-tuned deep neural network models. In regard to the encoding we were able to demonstrate a significant (>76%) radar data rate reduction while retaining most of the accuracy (<2% accuracy loss). Finally, in the early sensory fusion of the two modalities (binarized ADC radar data and DVS events), we were able to confirm higher accuracy than either modality alone. To the best of our knowledge, this is the first demonstration of sensory fusion of radar and DVS camera data aimed at classifying human gestures using only binary input information from the sensors.

Acknowledgment

This project has received funding from the ECSEL Joint Undertaking (JU) under grant agreement No. 826610, by the EU H2020 project MeMScales (Grant Agreement 871371), by the EUREKA cluster PENTA funded by Dutch authorities (PENTA2018e-17004-SunRISE), and by the Dutch Efficient Deep Learning project (grant P16-25-P7).

References

[1] A. Shaaban, W. Furtner, R. Weigel, and F. Lurz, "Evaluation of spiking neural networks for time domain-based radar hand gesture recognition," in *2022 23rd International Radar Symposium (IRS)*. IEEE, 2022, pp. 474–479.

[2] M. Arsalan, A. Santra, and V. Issakov, "Spiking neural network-based radar gesture recognition system using raw adc data," *IEEE Sensors Letters*, vol. 6, no. 6, pp. 1–4, 2022.

[3] I. J. Tsang, F. Corradi, M. Sifalakis, W. Van Leekwijck, and S. Latré, "Radar-based hand gesture recognition using spiking neural networks," *Electronics*, vol. 10, no. 12, p. 1405, 2021.

[4] F. Kreutz, P. Gerhards, B. Vogginger, K. Knobloch, and C. G. Mayr, "Applied spiking neural networks for radar-based gesture recognition," in *2021 7th International Conference on Event-Based Control, Communication, and Signal Processing (EBCCSP)*, 2021, pp. 1–4.

[5] A. Safa, F. Corradi, L. Keuninckx, I. Ocket, A. Bourdoux, F. Catthoor, and G. G. Gielen, "Improving the accuracy of spiking neural networks for radar gesture recognition through preprocessing," *IEEE Transactions on Neural Networks and Learning Systems*, 2021.

[6] M. Chmurski, G. Mauro, A. Santra, M. Zubert, and G. Dagasan, "Highly-optimized radar-based gesture recognition system with depthwise expansion module," *Sensors*, vol. 21, no. 21, 2021. [Online]. Available: https://www.mdpi.com/1424-8220/21/21/7298

[7] J. Stuijt, M. Sifalakis, A. Yousefzadeh, and F. Corradi, "μbrain: An event-driven and fully synthesizable architecture for spiking neural networks," *Frontiers in neuroscience*, vol. 15, p. 538, 2021.

[8] Y. Zhao, T. Liu, X. Feng, Z. Zhao, W. Cui, and Y. Fan, "New application: A hand air writing system based on radar dual view sequential feature fusion idea," *Remote Sensing*, vol. 14, no. 20, 2022. [Online]. Available: https://www.mdpi.com/2072-4292/14/20/5177

[9] C. Li, E. Tanghe, J. Fontaine, L. Martens, J. Romme, G. Singh, E. De Poorter, and W. Joseph, "Multistatic uwb radar-based passive human tracking using cots devices," *IEEE Antennas and Wireless Propagation Letters*, vol. 21, no. 4, pp. 695–699, 2022.

[10] M. González-Álvarez, J. Dupeyroux, F. Corradi, and G. C. De Croon, "Evolved neuromorphic radar-based altitude controller for an autonomous open-source blimp," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 85–90.

[11] X. Bai, Y. Hui, L. Wang, and F. Zhou, "Radar-based human gait recognition using dual-channel deep convolutional neural network," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 12, pp. 9767–9778, 2019.

[12] J. López-Randulfe, T. Duswald, Z. Bing, and A. Knoll, "Spiking neural network for fourier transform and object detection for automotive radar," *Frontiers in Neurorobotics*, vol. 15, 2021. [Online]. Available: https://www.frontiersin.org/article/10.3389/fnbot.2021.688344

[13] K. Seshadri, B. Akin, J. Laudon, R. Narayanaswami, and A. Yazdanbakhsh, "An evaluation of edge tpu accelerators for convolutional neural networks," 2021. [Online]. Available: https://arxiv.org/abs/2102.10423

[14] A. Kurniawan, *Introduction to NVIDIA Jetson Nano*. Berkeley, CA: Apress, 2021, pp. 1–6. [Online]. Available: https://doi.org/10.1007/978-1-4842-6452-2_1

[15] A. Yousefzadeh, G.-J. van Schaik, M. Tahghighi, P. Detterer, S. Traferro, M. Hijdra, J. Stuijt, F. Corradi, M. Sifalakis, and M. Konijnenburg, "Seneca: Scalable energy-efficient neuromorphic computer architecture," in *2022 IEEE 4th International Conference on Artificial Intelligence Circuits and Systems (AICAS)*. IEEE, 2022, pp. 371–374.

[16] M. Davies, N. Srinivasa, T.-H. Lin, G. Chinya, Y. Cao, S. H. Choday, G. Dimou, P. Joshi, N. Imam, S. Jain *et al.*, "Loihi: A neuromorphic manycore processor with on-chip learning," *Ieee Micro*, vol. 38, no. 1, pp. 82–99, 2018.

[17] S. Furber and P. Bogdan, *SpiNNaker-A Spiking Neural Network Architecture*. Now publishers, 2020.

[18] C. Brandli, R. Berner, M. Yang, S.-C. Liu, and T. Delbruck, "A 240x 180 130 db 3 μs latency global shutter spatiotemporal vision sensor," *IEEE Journal of Solid-State Circuits*, vol. 49, no. 10, pp. 2333–2341, 2014.

[19] Y.-H. Liu, S. Sheelavant, M. Mercuri, P. Mateman, and M. Babaie, "An ultralow power burst-chirp uwb radar transceiver for indoor vital signs and occupancy sensing in 40-nm cmos," *IEEE Solid-State Circuits Letters*, vol. 2, no. 11, pp. 256–259, 2019.

[20] G. Gallego, T. Delbrück, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A. J. Davison, J. Conradt, K. Daniilidis *et al.*, "Event-based vision: A survey," *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 1, pp. 154–180, 2020.

[21] A. Amir, B. Taba, D. Berg, T. Melano, J. McKinstry, C. Di Nolfo, T. Nayak, A. Andreopoulos, G. Garreau, M. Mendoza *et al.*, "A low power, fully event-based gesture recognition system," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 7243–7252.

[22] Y. Bi, A. Chadha, A. Abbas, , E. Bourtsoulatze, and Y. Andreopoulos, "Graph-based object classification for neuromorphic vision sensing," in *2019 IEEE International Conference on Computer Vision (ICCV)*. IEEE, 2019.

[23] J. Lien, N. Gillian, M. E. Karagozler, P. Amihood, C. Schwesig, E. Olson, H. Raja, and I. Poupyrev, "Soli: Ubiquitous gesture sensing with millimeter wave radar," *ACM Transactions on Graphics (TOG)*, vol. 35, no. 4, pp. 1–19, 2016.

[24] M. Ritchie, R. Capraru, and F. Fioranelli, "Dop-net: a micro-doppler radar data challenge," *Electronics Letters*, vol. 56, no. 11, pp. 568–570, 2020.

[25] J. López-Randulfe, N. Reeb, N. Karimi, C. Liu, H. Gonzalez, R. Dietrich, B. Vogginger, C. Mayr, and A. Knoll, "Time-coded spiking fourier transform in neuromorphic hardware," *IEEE Transactions on Computers*, 2022.