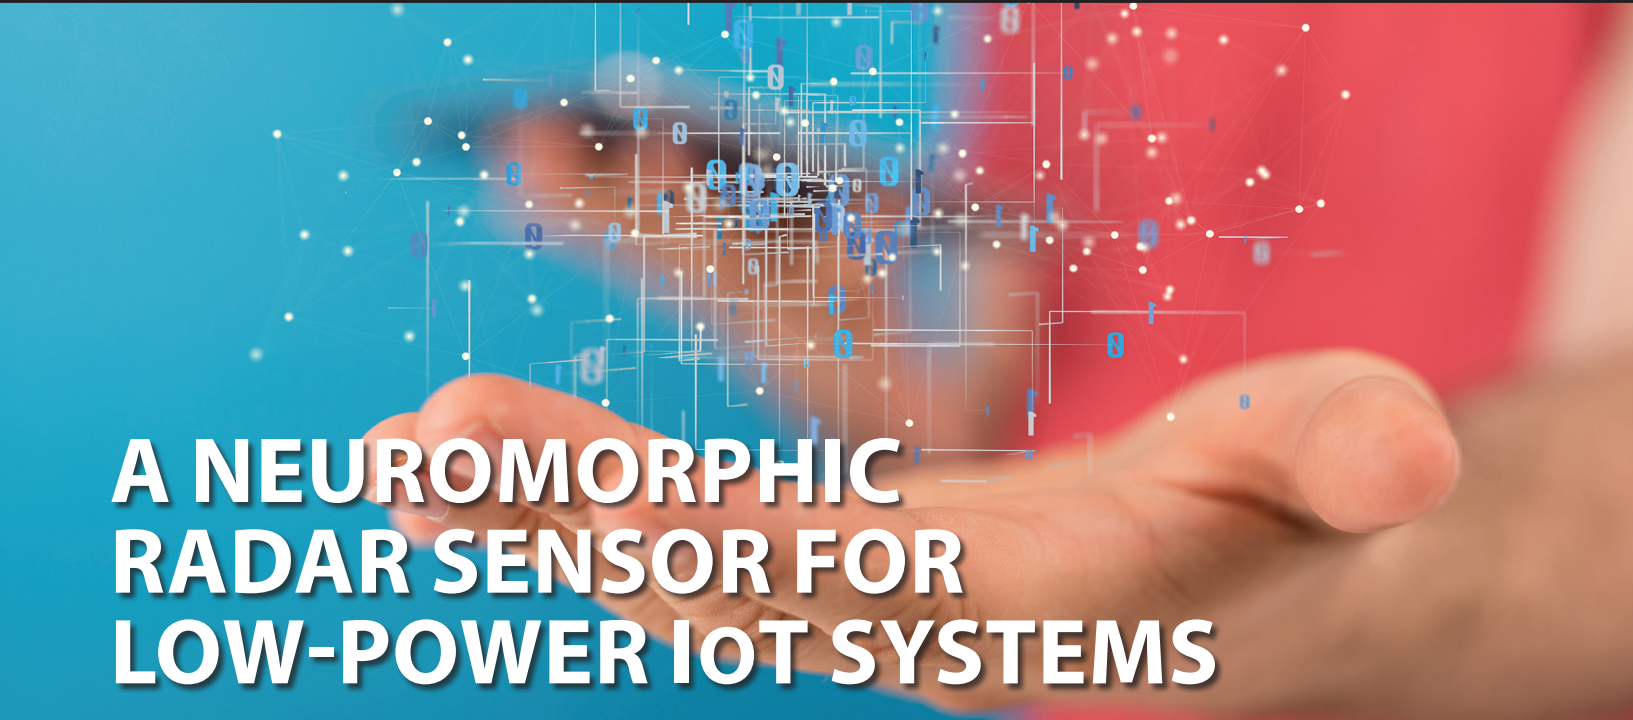


Kai Zheng University of California, San Diego, CA **Kun Qian** University of Virginia, Charlottesville, VA
Timothy Woodford MIT Lincoln Lab, Lexington, MA **Xinyu Zhang** University of California, San Diego, CA



A NEUROMORPHIC RADAR SENSOR FOR LOW-POWER IOT SYSTEMS

Excerpted from "NeuroRadar: A Neuromorphic Radar Sensor for Low-Power IoT Systems," from SenSys '23: *Proceedings of the 21st ACM Conference on Embedded Networked Sensor Systems*, with permission. <https://doi.org/10.1145/3625687.3625788> ©ACM 2023

As radar sensors become an integral component of Internet of Things (IoT) systems, the challenge of high power consumption poses a significant barrier, especially for battery-operated devices. This article introduces *NeuroRadar*, a groundbreaking solution that leverages a radar front-end capable of generating spike sequences, which can be efficiently processed by energy-saving Spiking Neural Networks (SNNs). We explore the innovative design and implementation of NeuroRadar, showcasing its effectiveness in applications like gesture recognition and human tracking. By achieving dramatically lower power consumption compared to traditional radar systems, NeuroRadar represents a new paradigm in energy-efficient IoT sensing.

The Internet of Things (IoT) is rapidly transforming modern technology, with applications spanning from smart homes and healthcare monitoring to industrial automation and intelligent transportation systems. Central to many of these applications are radar sensors, which provide essential functions such as motion detection, gesture recognition, and activity classification. However, the high power consumption of radar hardware presents a significant challenge, especially for battery-operated IoT devices and wearables, where energy efficiency and battery life are para-

mount. This challenge is further exacerbated by the reliance on power-intensive artificial neural networks (ANNs) for signal processing in many smart sensing applications.

Recent advancements in neuromorphic engineering have led to the development of Spiking Neural Networks (SNNs)[1] and dedicated neuromorphic circuits[2], which more closely emulate the efficiency of sensory signal processing found in the brain. SNNs are designed to mirror the pulse-based behavior of the human nervous system, consisting of spiking

neurons and the synaptic connections between them. When realized on dedicated neuromorphic circuits, SNNs demonstrate exceptional energy efficiency, surpassing traditional von Neumann computing units by orders of magnitude[3]. This revolution in neuromorphic computing has also spurred the development of cutting-edge neuromorphic sensing hardware, such as the energy-efficient, fast-response event camera[4].

In this article, we explore NeuroRadar[5], a novel low-power radar sensing system that leverages the power of neuromorphic sensing and computing. NeuroRadar draws inspiration from neuromorphic sensors that mimic mammalian sensory systems, generating event-triggered outputs in response to external stimuli, as illustrated in Figure 1. Unlike traditional radars that produce continuous frame-based outputs, NeuroRadar generates spiking patterns upon detecting motion in its environment. While other so-called "SNN radars"[6,7] continue to rely on traditional CPUs or digital signal processing (DSP) units, NeuroRadar adopts a fully neuromorphic architecture, processing all information in the spike domain.

SYSTEM OVERVIEW

NeuroRadar comprises three main components: a sensor front-end, spike encoders, and spike processors (Figure 1). The sensor front-end detects ambient motion, and the spike encoders convert these signals into spike sequences, known as spike trains. These spike trains are then processed directly by energy-efficient SNNs.

Sensor front-end

NeuroRadar employs a drastically simplified RF front-end that eliminates most power-intensive active RF components found in traditional radars, retaining only a low-power free-running oscillator. NeuroRadar detects environmental changes using the self-injection locking (SIL) principle[8], where the oscillator's frequency is modulated by motion in the surrounding area. By demodulating this frequency shift, the system generates a baseband signal carrying motion information. Despite its simplicity, NeuroRadar maintains a reasonable level of sensitivity due to the inherent properties of the SIL architecture, which enhance signal strength. Notably, the power consumption of each RF front-end channel is just 240 μ W in an integrated circuit implementation.

However, a single SIL sensor cannot provide angular resolution and accurate range information, as it only senses environmental motion. To overcome this limitation, NeuroRadar draws inspiration from the compound eyes found in certain biological organisms[9], proposing an array of SIL sensors with strategically selected carrier frequencies. This design allows the system to infer target direction by exploiting the phase difference across the sensors and resolves range ambiguity using frequency diversity. The original NeuroRadar paper[5] includes an optimization framework that determines the number of sensors and their frequency permutation for accurate target localization.

Spike encoder

To enable end-to-end SNN signal processing, NeuroRadar uses an analog spike encoding circuit that directly transforms baseband signals into spike trains. The spike encoder




FIGURE 1. Analogy between biological sensory systems and the NeuroRadar sensor.

NeuroRadar achieves energy-efficient sensing by emulating the structure and functionality of these biological systems.

is based on a biological neuron model[10], preserving all essential sensing information within the spike sequences. The encoder employs spike rate encoding, where the firing frequency increases linearly with the input signal amplitude. These spike encoding circuits operate entirely in an event-driven manner, generating spikes only when motion is detected by the sensor front-end, as shown in Figure 2. This ensures that the system remains idle in the absence of motion, further enhancing energy efficiency. The spike sequences are then processed directly by SNNs on neuromorphic computing systems, eliminating the need for any non-spike-based computing units. Figure 3 illustrates the implementation of a single-channel SIL radar.

Spike processor

Once motion signals are converted into multiple parallel spike trains, SNNs are designed to process these spiking signals and extract spatiotemporal features. The SNN structure includes three main components: spike buffering units, convolution layers, and spike decoders (Figure 4).

- **Spike Buffering Units:** These units consist of cascaded time delay units, each imposing a consistent delay of a certain number of clock ticks. The output spikes then progress to the next-stage delay unit.




FIGURE 2. Spike encoding circuit. The signal from each SIL channel is encoded into two spike trains, representing the positive and negative parts respectively.

Once the input sequence is completed, the spike buffering units concatenate the outputs from all delay units, presenting the spikes concurrently to the subsequent layer.

- **Convolution Layers:** After flattening the temporal dimension of the input spike sequences through the spike buffering units, convolution layers extract the spatiotemporal features of the spike sequences.
- **Spike Decoders:** SNNs are trained so that output values are represented by the spike firing rate of neurons in the final layer. Spike decoders, which are essentially low-pass filters, convert the spike density into continuous values that can be interpreted by the sensing applications.

This comprehensive SNN processing workflow enables NeuroRadar to deliver application-specific sensing results with superior energy efficiency.

CASE STUDY

To demonstrate the capabilities of NeuroRadar, we conducted two case studies: hand gesture recognition and human tracking. For each case, we collected multi-channel spike data and trained and tested the SNNs. The SNN simulations were performed on the Intel Loihi neuromorphic processor[2] using the NengoDL framework[11]. In the Loihi implementation, the SNN operates on discretized time steps, each corresponding to 1 ms.

Gesture Recognition

We first illustrate NeuroRadar’s ability to perform hand gesture recognition. In this case study, we defined a gesture set comprising 12 different gestures, as shown in Figure 5. The gesture set includes diverse hand movements within a 3D space (e.g., push, pull, left, right, up, down), with some gestures requiring simultaneous movement of both hands.

A three-channel NeuroRadar setup was employed. Two antennas were placed on a horizontal line, while the third antenna was positioned above the horizontal line, forming an equilateral triangle with the other two antennas. Gestures were performed in front of the antenna plane, facing the center of the triangle. The spacing between antennas was designed to be comparable to the displacement of a hand during gestures, resulting in more




FIGURE 3. Single-channel SIL radar system. The ambient motion modulates the oscillating frequency of the SIL oscillator, which is then demodulated using a delay and mixing circuit. The resulting signals are amplified and converted into spike trains by the amplifiers and spike encoders.

distinguishable signal patterns. The elevated antenna provided richer information for vertical hand movements, such as “swipe up” or “swipe down.”

A total of 2,400 samples were collected, with 200 samples for each gesture. Each gesture sample contained sequences of spikes over a 1.5-second duration. The samples were divided into a training set (1,920 samples) and a test set (480 samples) using a random 80/20 split.

NeuroRadar’s SNN achieved an accuracy of 94.6% after running for 80 time steps. This indicates that after the completion of a gesture input, NeuroRadar requires only 80 ms to produce a reliable result, which is sufficient for most applications. Table 1 summarizes the gesture recognition outcomes.

The power consumption of the RF front-end and the SNN is merely 830 μ W and

65 μ W, respectively. Compared to other SNN-based gesture recognition systems[6,7] that still rely on traditional computing units (e.g., CPU or DSP) for radar signal pre-processing, NeuroRadar achieves a power consumption reduction between 78% and 93% in terms of signal processing (pre-processing and SNN).

Human Tracking

The second case study focuses on localizing a single moving human in an indoor environment. NeuroRadar employs a 6-sensor array with diverse carrier frequencies to achieve this. Since NeuroRadar detects only moving targets, a volunteer was asked to walk randomly within the radar’s field of view. The maximum distance of the target from the radar was approximately 6 meters, with a viewing angle of about 90 degrees. Ground truth for location and speed was obtained using a commercial depth camera (ZED-2i). We collected 3,600 seconds of continuous data, which was then split 80/20 into training and testing datasets. The continuous data was further segmented into 2-second short frames with a 75% overlap, resulting in each short frame becoming a training or test sample. This segmentation yielded a total of 5,742 training samples and 1,422 test samples.

Figure 6 displays the localization results by combining the output from consecutive frames. Similar to the gesture recognition use case, the SNN requires enough timesteps to produce a reliable result. The results show that after approximately 150 timesteps, a localization accuracy of 1 meter can be achieved. The mean squared error for speed estimation stabilizes at 0.25 m²/s². This indicates a tracking delay of 150 ms,




FIGURE 4. SNN processing pipeline. All the processing is conducted in the spike domain.




FIGURE 5. Gesture set definition. Gestures (1)-(8) are single-hand; (9)-(12) are double-hand.

TABLE 1. Confusion matrix for gesture recognition

	1	2	3	4	5	6	7	8	9	10	11	12
1	0.94								0.03			0.03
2		0.95										0.05
3			0.93					0.02	0.05			
4		0.03		0.91				0.03				0.03
5					0.97				0.03			
6		0.04				0.96						
7							1.00					
8	0.02							0.98				
9	0.02		0.04						0.92			0.02
10	0.02	0.07		0.02		0.03		0.03		0.83		
11											0.98	0.02
12							0.02					0.98

which is adequate for low-velocity indoor applications. However, because the system filters spike sequences to achieve continuous values, some errors are inevitable, impacting overall accuracy.

Thanks to its simple SIL structure and power-efficient design, NeuroRadar consumes only 2.03 mW of power (1.44 mW for the front-end and 0.59 mW for the SNN). Compared to Doorpler [12], another motion-based surveillance radar system, NeuroRadar achieves a 1-2 orders of magnitude reduction in front-end power consumption. The combination of the SIL array and neural network allows NeuroRadar to obtain richer and more accurate sensing information. Unlike Doorpler, which merely detects crossing events and their direction, NeuroRadar offers both location and speed estimation. Additionally, SNN processing significantly reduces computational power, resulting in an overall system power consumption reduction of 97%.

SUMMARY

In this article, we reviewed NeuroRadar, a novel and pioneering approach in radar systems that fully embraces the principles of neuromorphic sensing. By jointly designing analog hardware and spike-based signal processing, NeuroRadar achieves exceptional energy efficiency. Through case studies on gesture recognition and human localization, NeuroRadar has demonstrated its capabilities while maintaining significantly lower power consumption compared to traditional radar systems. This research represents a significant advancement, offering a unique and innovative solution for radar sensing in energy-constrained IoT devices. ■

Acknowledgements

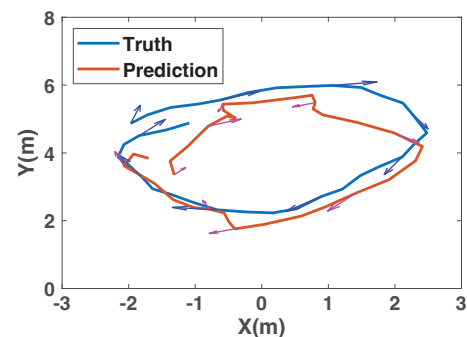
The work reported in this article is supported in part by the NSF under Grants CNS-2312715, CNS-1901048, CNS-1925767, and CNS-2128588.

Kai Zheng is currently a PhD Candidate at University of California, San Diego. He received his BS from Fudan University in 2015 and MS from New York University in 2019. His research focuses on mmWave sensing and low-power IoT systems. He is the recipient of SenSys 2023 Best Paper Award.

Kun Qian received his PhD from Tsinghua University in 2019. Prior to joining the University of Virginia in 2024, he was a Postdoctoral

REFERENCES

- [1] W. Maass. 1997. Networks of spiking neurons: The third generation of neural network models. *Neural Networks*, vol. 10, no. 9, 1659–1671.
- [2] M. Davies, N. Srinivasa, T.-H. Lin, G. Chinya, Y. Cao, S. H. Choday, G. Dimou, P. Joshi, N. Imam, S. Jain et al. 2018. Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro*, vol. 38, no. 1, 82–99.
- [3] P. Blouw, X. Choo, E. Hunsberger, and C. Eliasmith. 2019. Benchmarking keyword spotting eZiciency on neuromorphic hardware. *Proceedings of the 7th Annual Neuroinspired Computational Elements Workshop*, 1–8.
- [4] P. Lichtsteiner, C. Posch, and T. Delbruck. 2008. A 128×128 120 db 15 μs latency asynchronous temporal contrast vision sensor. *IEEE Journal of Solid-State Circuits*, vol. 43, no. 2, 566–576.
- [5] K. Zheng, K. Qian, T. Woodford, and X. Zhang. 2023. NeuroRadar: A neuromorphic radar sensor for low-power IoT systems. *Proceedings of the 21st ACM Conference on Embedded Networked Sensor Systems*, 223–236.
- [6] M. Arsalan, A. Santra, and V. Issakov. 2022. Radarsnn: A resource eZicient gesture sensing system based on mm-wave radar. *IEEE Transactions on Microwave Theory and Techniques*, vol. 70, no. 4, 2451–2461.
- [7] M. Arsalan, A. Santra, and V. Issakov. 2023. Power-efficient gesture sensing for edge devices: Mimicking fourier transforms with spiking neural networks. *Applied Intelligence*, vol. 53, no. 12, 15 147–15 162.
- [8] F.-K. Wang, C.-J. Li, C.-H. Hsiao, T.-S. Horng, J. Lin, K.-C. Peng, J.-K. Jau, J.-Y. Li, and C.-C. Chen. 2010. A novel vital-sign sensor based on a self-injection-locked oscillator. *IEEE Transactions on Microwave Theory and Techniques*, vol. 58, no. 12, 4112–4120.
- [9] M. Henning, G. Rasmus-Traslosheros, B. Bär and M. Silies. 2022. Populations of local direction-selective cells encode global motion patterns generated by self-motion. *Science Advances*, vol. 8, no. 31, eabi7112.
- [10] W. Gerstner, W. M. Kistler, R. Naud, and L. Paninski. 2014. *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press.
- [11] D. Rasmussen. 2019. NengoDL: Combining deep learning and neuromorphic modelling methods. *Neuroinformatics*, vol. 17, no. 4, 611–628.
- [12] A. Kalyanaraman, E. Soltanaghaei, and K. Whitehouse. 2019. Doorpler: A radar-based system for real-time, low power zone occupancy sensing. 2019. *Proceedings of IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*. 42–53.

**FIGURE 6.** NeuroRadar localization result.

researcher at the University of California, San Diego. His research interests include wireless sensing and ubiquitous computing. He received an ACM MobiCom Best Paper Award in 2020 and an ACM SenSys Best Paper Award in 2023.

Timothy Woodford holds a PhD in Electrical and Computer Engineering from University of California, San Diego. He is now a Research Scientist at MIT Lincoln Lab. His research focuses on mmWave sensing and communication.

Xinyu Zhang is an associate professor at University of California, San Diego. He received his PhD in Computer Science and Engineering from the University of Michigan. His research area includes wireless networking and ubiquitous computing. He received ACM's Best Paper Award twice in top ACM conferences and the ACM Rockstar Award in 2023.