



King's Research Portal

DOI:

[10.1109/TCCN.2023.3236940](https://doi.org/10.1109/TCCN.2023.3236940)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Chen, J., Skatchkovsky, N., & Simeone, O. (2023). Neuromorphic Wireless Cognition: Event-Driven Semantic Communications for Remote Inference. *IEEE Transactions on Cognitive Communications and Networking*, 9(2), 252-265. <https://doi.org/10.1109/TCCN.2023.3236940>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Neuromorphic Wireless Cognition: Event-Driven Semantic Communications for Remote Inference

Jiechen Chen, *Student Member, IEEE*, Nicolas Skatchkovsky, *Member, IEEE*,
Osvaldo Simeone, *Fellow, IEEE*

Abstract

Neuromorphic computing is an emerging computing paradigm that moves away from batched processing towards the online, event-driven, processing of streaming data. Neuromorphic chips, when coupled with spike-based sensors, can inherently adapt to the “semantics” of the data distribution by consuming energy only when relevant events are recorded in the timing of spikes and by proving a low-latency response to changing conditions in the environment. This paper proposes an end-to-end design for a neuromorphic wireless Internet-of-Things system that integrates spike-based sensing, processing, and communication. In the proposed NeuroComm system, each sensing device is equipped with a neuromorphic sensor, a spiking neural network (SNN), and an impulse radio (IR) transmitter with multiple antennas. Transmission takes place over a shared fading channel to a receiver equipped with a multi-antenna impulse radio receiver and with an SNN. In order to enable adaptation of the receiver to the fading channel conditions, we introduce a hypernetwork to control the weights of the decoding SNN using pilots. Pilots, encoding SNNs, decoding SNN, and hypernetwork are jointly trained across multiple channel realizations. The proposed system is shown to significantly improve over conventional frame-based digital solutions, as well as over alternative non-adaptive training methods, in terms of time-to-accuracy and energy consumption metrics.

Index Terms

Neuromorphic computing, spiking neural networks, semantic communications.

The work of Osvaldo Simeone and Nicolas Skatchkovsky was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme, grant agreement No. 725731, by an Open Fellowship of the EPSRC with reference EP/W024101/1, and by the European Union through project CENTRIC (101096379), and the work by Jiechen Chen was funded by the China Scholarship Council and King’s College London for their Joint Full-Scholarship (K-CSC) under Grant CSC202108440223.

The authors are with the King’s Communications, Learning and Information Processing (KCLIP) lab, King’s College London, London, WC2R 2LS, UK. (email:{jiechen.chen, nicolas.skatchkovsky, osvaldo.simeone}@kcl.ac.uk).

I. INTRODUCTION

A. Context and Motivation

The recent rollout of 5G around the world has marked the start of a switch of telecom systems from *network-centric carriers of bits* to *user-centric distributed processors of intelligence*. A key element of this switch will be the integration of wireless systems with sensing and cognition, a technology trend we will refer to as *wireless cognition*. In this context, this paper is motivated by the two following paradigm shifts that are widely envisioned as central to 6G:

- *From universality to goal-driven specialization*: With the emergence of wireless cognition, there is a need to develop semantics-aware solutions that integrate sensing, communication, and computing, tailoring resource consumption to the goals of the task at hand [1–5].
- *From hardware agnostic to hardware-constrained design*: By assuming universal computing architectures, the conventional design of communication systems is agnostic to the specific hardware systems deployed at the communicating nodes. This approach fails to acknowledge the critical importance of the computing architecture in the effective and efficient implementation of semantic tasks (see, e.g., [6]).

Neuromorphic sensing and computing are emerging as alternative, *brain-inspired*, paradigms for efficient data collection and semantic signal processing. The main features of the technology are energy efficiency, native event-driven processing of time-varying semantic sources, spike-based computing, and always-on on-hardware adaptation [7, 8].

- *Neuromorphic sensors* encode information in the timing of spikes, and include neuromorphic cameras, silicon cochleas, and brain-computer interfaces. As a general principle of operation, spikes are produced only when relevant changes occur in the signals being sensed [9–12].
- *Neuromorphic processors*, also known as *spiking neural networks (SNNs)*, are networks of dynamic spiking neurons that mimic the operation of biological neurons [13]. Spiking neurons communicate and process with the timings of spikes [14]. When implemented on specialized – digital or mixed analog-digital – hardware or on tailored FPGA configurations, SNNs have minimal idle and operating energy cost, and consume as little as *a few picojoules* per spike [15].

Current commercial use cases of neuromorphic technologies range from drone monitoring via Dynamic Vision Sensor (DVS) cameras [16, 17] through the development of brain-computer

interfaces¹ to the development of fast and accurate COVID-19 antibody testing². Neuromorphic computing platforms include Intel’s Loihi SNN chip, IBM’s TrueNorth, and Brainchip’s Akida [18, 19]. This work views the emergence of neuromorphic technologies as a unique opportunity for the development of *goal-driven, specialized, and hardware-constrained wireless cognition*.

To date, work on the integration of wireless connectivity and neuromorphic systems has been very limited, including only a specific implementation introduced for biomedical applications in [20].

B. NeuroComm

This paper introduces *NeuroComm*, a wireless cognition system that integrates neuromorphic sensing, computing, and spike-based communications. As illustrated in Fig 1, NeuroComm consists of a wireless Internet-of-Things (IoT) network in which multiple wireless sensing devices communicate over a shared wireless link to an edge device. We target an *always-on* monitoring application in which the environment being observed by the sensors presents *sparse* activity, and *inference* at the receiving edge device should be carried out in a *timely* manner.

An example of the scenarios under consideration is given by distributed sensors deployed to monitor the presence of drone activity, e.g., in an airport. In such applications, sensors capture uninformative signals most of the time, while needing to promptly enable object recognition and detection at the receiving processor when relevant activity is present in the system.

As depicted in Fig 2-(a), a conventional deployment for remote inference would rely on digital *frame-based* sensing, processing, and computing blocks. This approach has two major drawbacks: (i) *Semantic-agnostic energy consumption*: Conventional digital sensing devices would generally sense and transmit frames continuously, irrespective of the level of “information” present in the sensed signal for the given application; (ii) *Latency quantization*: Sensed information would be packaged into frames, enabling the receiver to enhance its inference accuracy only upon the reception and processing of entire frames.

In NeuroComm, as illustrated in Fig 2-(b), we propose to replace the digital frame-based blocks of the conventional deployment in Fig 2-(a) with neuromorphic sensing, computing, and transmission blocks. The main goals of these system-level innovations are as follows: (i)

¹<https://neuralink.com/>

²<https://bit.ly/3r0NeoS>

Semantic-aware energy consumption: First, we wish to ensure an energy consumption level throughout the sensing-processing-communication chain that adapts to the dynamics of the environment being monitored; *(ii) Enhanced time-to-efficiency:* Second, we aim at enabling a graceful improvement in the inference performance over time as more information on the environment is sensed and communicated.

These goals are met by leveraging the *event-driven* nature of neuromorphic sensing and computing, as well as the synergy between spike processing and pulse-based transmission via *impulse radio* [21]. Neuromorphic sensors, SNNs, and impulse radio blocks consume energy only when spikes are produced, reflecting patterns of activity in the monitored scene. The use of impulse radio is also appealing in light of the role of wideband transmission for THz communications in 6G, as well as in existing low-power communication standards like IEEE 802.15.4 [22, 23].

C. Main Contributions

As depicted in Fig. 3, NeuroComm implements a novel form of *joint source-channel coding* whereby data from neuromorphic sensors are pre-processed by SNNs and directly mapped into electromagnetic pulses on shared wireless channels in an online, streaming, fashion. The receiver processes the received signals via an SNN to carry out inference. This approach differs from the standard *separate source-channel coding* solution illustrated in Fig. 4 that is implemented by digital frame-based systems. In the conventional system, data from sensors are first compressed and then encoded on a frame-by-frame basis.

The main contributions of this paper are summarized as follows.

- We introduce NeuroComm, a novel architecture for wireless cognition integrating neuromorphic sensing, processing, and communications.
- We propose an end-to-end design of the NeuroComm architecture based on supervised learning via surrogate gradient descent methods.
- In order to enable adaptation to channel conditions, we introduce a hypernetwork-based approach. The hypernetwork is trained to take as input signals received from pilot symbols to control the weights applied by the receiving SNN.
- Extensive numerical results are provided that demonstrate the advantages of the proposed architecture and design over conventional frame-based digital solutions, as well as over alter-

native non-adaptive training methods, in terms of time-to-accuracy and energy consumption metrics.




Fig. 1. This paper focuses on an edge-based IoT system in which data are collected and pre-processed in a streaming fashion by mobile sensors so that inference can be carried out continuously at another edge device. The proposed NeuroComm solution integrates neuromorphic sensing, neuromorphic computing – via spiking neural networks (SNNs) – and impulse radio.

D. Related Work

To the best of our knowledge, the only paper that has proposed the idea of end-to-end learning for a communication link with neuromorphic encoder and decoder is the conference publication [24] by a subset of the authors of this paper. This is a preliminary study focusing on a special case of the NeuroComm system in which a single transmitter operates over a frequency-flat channel with input data collected with a neuromorphic camera. In contrast, the current work investigates the general case of multiple devices communicating over a shared fading channel.

A paper [25] by a team at Ericsson has considered the distributed implementation of an SNN on a packet-based wireless channel, in which the SNN is trained offline in a channel-agnostic manner. Moreover, a recent work [26] proposed a decision-feedback equalizer based on an SNN as a solution for an energy-efficient receiver.

There is a rich literature on joint source-channel coding for conventional artificial neural networks (ANNs). For example, [27–29] investigated wireless image transmission based on




Fig. 2. (a) Illustration of a conventional digital implementation of an IoT link with remote inference, in which sensed information is transformed into bits, packaged into frames, encoded, and transmitted; (b) Illustration of the proposed NeuroComm system in which digital blocks are replaced by neuromorphic blocks, including neuromorphic sensing and processing via SNNs. In NeuroComm, spikes are communicated by means of impulse radio (IR) and processed via SNNs.

convolutional neural networks. As another example, the authors in [30] optimized intelligent reflecting surfaces-assisted wireless communication systems in an end-to-end manner by treating the transmission chain as a joint source-channel coding problem.

As for the applications of hypernetworks in communications, reference [31] studied an application to MIMO detection under highly correlated channels to avoid retraining for each channel realization. The paper [32] proposed a hypernetwork-based architecture for recurrent neural networks that leverages partial uplink-downlink reciprocity and temporal correlation of the fading channel to perform channel estimation in frequency-division duplex systems. The use of hypernetworks coupled with SNNs appears to be novel.

E. Organization

The remainder of the paper is organized as follows. Section II presents the system model for multi-device remote inference. The conventional frame-based digital transmission is reviewed for reference in Section III. The NeuroComm architecture is proposed in Section IV, while Section V




Fig. 3. A block diagram illustrating the main operations involved in the implementation of the proposed NeuroComm system (IR stands for impulse radio).

presents the integration of NeuroComm with a hypernetwork and the resulting training problem. Experimental setting and results are described in Section VI. Finally, Section VII concludes the paper.

II. SYSTEM MODEL

In this paper, we study the wireless edge-based IoT system illustrated in Fig. 1, whose goal is to implement continuous remote inference via distributed sensors and an edge processor in an energy-efficient and timely manner. To this end, the system integrates neuromorphic computing via spiking neural networks (SNNs) and impulse radio (IR) transmission. In this section, we formally describe setting, channel model, the operation of neuromorphic sensors, and the transmission model.

A. Setting

As illustrated in Fig. 1, we consider a remote inference system with one receiver collecting information from K sensing devices. Each device $k \in \mathcal{K} = \{1, \dots, K\}$ is equipped with a neuromorphic sensor and with N^T antennas, while the receiver has N^R antennas. Communication takes place over a shared multi-path fading channel under average or peak power constraints.

A conventional digital sensor, such as a digital camera, would capture natural signals via sampling and analog-to-digital conversion. Accordingly, in a continuous, always-on, monitoring application, a digital sensor would produce a constant stream of bits, typically packaged into frames. In this paper, in order to leverage efficiencies related to the assumed sparsity of the dynamics of the environment being monitored, we assume that sensing devices are equipped with neuromorphic sensors, such as DVS cameras [33]. We recall that, by the principle of operation of a neuromorphic sensor, the presence of a spike generally indicates the occurrence of an event. For instance, with DVS cameras, a spike is produced as the result of a sufficiently large change in brightness at a pixel [17].

We adopt a discrete-time model, with discrete time index $l = 1, 2, \dots, L$ used to denote the time samples produced by the neuromorphic sensors. We assume the devices to be time synchronized. The sensor available at each device produces D_u spiking signals. For example, for a DVS camera, D_u is the number of pixels. Accordingly, the sensor at the k th device outputs a $D_u \times 1$ vector \mathbf{u}_l^k with binary entries at each discrete time l . Each binary entry $u_{l,m}^k$ of the $D_u \times 1$ vector $\mathbf{u}_l^k = [u_{l,1}^k, \dots, u_{l,D_u}^k] \in \{0, 1\}^{D_u}$ represents the presence ($u_{l,m}^k = 1$), or the absence ($u_{l,m}^k = 0$) of a spike in the m th spiking signal at time l . As mentioned, a spike may indicate that the corresponding m th pixel has undergone a sufficiently large change in luminance during the l th time interval.

The signals \mathbf{u}_l^k captured by each k th neuromorphic sensor are processed and transmitted over a fading multiple access channel to a receiver. The receiver's goal is to carry out an inference task, such as classification, based on the received signals. In the next two sections, we will first consider for reference a conventional digital communication and processing chain depicted in Fig. 4, and we will then introduce the proposed NeuroComm system, illustrated in Fig. 3.

B. Channel Model

We assume a general multi-path propagation model, which may account, e.g., for propagation in the sub-6GHz band, as in the IEEE 802.15.4z standard, or for millimeter wave and Terahertz (THz) bands [34]. Accordingly, the continuous-time channel response between the m th antenna of device k and the n th antenna of the receiver can be expressed as

$$h_{n,m}^k(t) = \sum_{p=1}^{N_P} \beta_{n,m,p}^k g(t - \tau_{n,m,p}^k), \quad (1)$$

where N_P is the total number of resolvable propagation paths, including possibly a line-of-sight (LoS) component; $\beta_{n,m,p}^k$ and $\tau_{n,m,p}^k$ represent the complex amplitude and delay of the p th path, respectively; and $g(\cdot)$ represents the matched filter response (see, e.g., [35]). The maximum delay of all channels is denoted as T_h , which is measured in seconds, i.e., we have $h_{n,m}^k(t) = 0$ for $t \geq T_h$ and for all values of n , m , and k . The channel is assumed to be constant for the duration of a transmission frame. A frame is assumed to be sufficiently long to encompass the presentation of input signals \mathbf{u}_l^k of L samples, $l = 1, \dots, L$, from the neuromorphic sensor to each device k .

C. Transmission Model

The modulated signal emitted by the m th antenna of the impulse radio transmitter k is denoted as $s_m^k(t)$ and is assumed to have bandwidth $1/T_c$, where T_c is measured in seconds. We write the overall duration of a transmission frame as LL_bT_c , where integer L_b is referred to as the *bandwidth expansion factor*. Since L is the number of samples produced by the sensors per frame, when sampled at Nyquist rate $1/T_c$, the transmitted and received signals present L_b samples for each sample of the sensed signals.

The signal obtained by the n th antenna of the receiver is given by the superposition of the signals transmitted by all the N^T antennas at the K devices as

$$y_n(t) = \sum_{k=1}^K \sum_{m=1}^{N^T} s_m^k(t) * h_{n,m}^k(t) + z_n(t), \quad (2)$$

where $z_n(t)$ is the additive white Gaussian noise with power spectral density N_0 ; “*” denotes the convolution operation; and $h_{n,m}^k(t)$ is the (complex) channel input response in (1) between antenna m of device k and receive antenna n .

Sampling signal (2) at rate $1/T_c$ yields the discrete-time signal $y_{i,n} = y_n(iT_c)$, for $i = 1, 2, \dots, LL_b$, across the time frame. Having a maximum delay spread of T_h , upon sampling, the effective discrete-time channels have $L_h = T_h/T_c$ taps. Accordingly, the channel between antenna m of device k and antenna n of the receiver is described by the $L_h \times 1$ vector $\mathbf{h}_{n,m}^k = [h_{n,m}^k(0), h_{n,m}^k(T_c), \dots, h_{n,m}^k((L_h - 1)T_c)]^T$. Therefore, the received discrete-time signal $y_{i,n}$ at the n th receive antenna can be expressed as

$$y_{i,n} = \sum_{k=1}^K \sum_{m=1}^{N^T} (\mathbf{h}_{n,m}^k)^T \mathbf{s}_{i,m}^k + z_{i,n}, \quad (3)$$

where we have defined the $L_h \times 1$ vector $\mathbf{s}_{i,m}^k = [s_{i,m}^k, s_{i-1,m}^k, \dots, s_{i-L_h+1,m}^k]^T$ with $s_{i,m}^k = s_m^k(iT_c)$, and $s_{i,m}^k = 0$ if $i \leq 0$; and $z_{i,n} = z_n(iT_c) \sim \mathcal{CN}(0, N_0)$.

We consider two different types of constraints on the transmitted signals under the assumption that the transmitters are not aware of the current channel realization.

- *Per-frame energy constraint:* Collecting the transmitted samples from all the N^T antennas of device k in a frame in the $LL_b N^T \times 1$ vector \mathbf{s}^k , the per-frame energy constraint imposes the inequality as

$$\|\mathbf{s}^k\|^2 \leq E_{\text{fr}}^k, \quad (4)$$

where E_{fr}^k is the maximal per frame transmission energy for transmitter k .

- *Per-symbol energy constraint:* The per-frame energy constraint (4) is well suited for conventional frame-based system, and is not a natural choice for systems operating in a streaming fashion. Therefore, we also consider the per-symbol energy constraint as

$$|s_{i,m}^k|^2 \leq E_s^k, \quad (5)$$

where E_s^k is the per-symbol energy constraint for transmitter k .

D. Pilot Transmission

In order to enable adaptation of the receiver to changing channel conditions, we assume that pilot symbols are transmitted by each device prior to the start of data transmission in each frame. Pilot transmissions from different devices are assumed to be orthogonal. Accordingly, device k transmits L_p symbols at rate $1/T_c$ from each antenna. Writing as $s_{i,m}^k$ with $i = -L_p + 1, -L_p + 2, \dots, 0$, as the pilot symbols sent by device k from antenna m , the signal received from device k during the training phase can be written as

$$y_{i,n}^k = \sum_{m=1}^{N^T} (\mathbf{h}_{n,m}^k)^T \mathbf{s}_{i,m}^k + z_{i,n}, \quad (6)$$

where we have defined the $L_h \times 1$ vector $\mathbf{s}_{i,m}^k = [s_{-L_p+1+i,m}^k, s_{-L_p+i,m}^k, \dots, s_{-L_p-L_h+i+2,m}^k]^T$, and $s_{i,m}^k = 0$ if $i \leq -L_p$. We impose the energy constraint

$$\|\mathbf{s}_p^k\|^2 \leq E_p^k, \quad (7)$$

which applies to the vector \mathbf{s}_p^k of all pilot symbols $s_{i,m}^k$ transmitted at time instants $i = -L_p + 1, \dots, 0$ from all antennas $m = 1, \dots, N^T$, and E_p^k is the maximum pilot energy constraint for transmitter k .




Fig. 4. Block diagram illustrating the main operations involved in a standard digital frame-based wireless remote inference system with neuromorphic sensing.

III. CONVENTIONAL FRAME-BASED DIGITAL TRANSMISSION

For reference, in this section we review the conventional frame-based digital system illustrated in Fig. 4. As discussed in Sec. I, this solution implements *separate source-channel coding* by processing the signals produced by the sensors via frame-based source and channel encoding and decoding blocks at the transmitters and at the receiver. While the sensors are neuromorphic for consistency with the setting under study in this work, unlike in the proposed NeuroComm architecture, all other blocks adopt standard digital computing and communication technologies.

To elaborate, as shown in Fig. 4, at each device k , the samples \mathbf{u}^k produced by the neuromorphic sensor are first divided into *subframes*. Each subframe is compressed via a source encoder, and then channel encoded. The channel encoded sequence is transmitted using standard digital modulation on the shared channel to the receiver. The receiver carries out channel decoding for the signals sent by all devices. The decoded signals are decompressed via the source decoder to recover an estimate of the original subframes produced by the sensors. These estimates are demultiplexed to produce an estimate of the sensed signals \mathbf{u}^k with $k = 1, \dots, K$, and then used for inference. Inference is carried out by a standard ANN at the receiver.

A. Separate Source-Channel Coding

The samples \mathbf{u}_l^k produced by each k th neuromorphic sensor across time instant $l = 1, \dots, L$ are formatted into subframes of L_{enc} samples each. As a result, $F = \lceil L/L_{\text{enc}} \rceil$ subframes \mathbf{u}_f^k

with $f = 1, \dots, F$ are produced, with each subframe \mathbf{u}_f^k containing $D_u L_{\text{enc}}$ binary digits. Note that if the ratio L/L_{enc} is not an integer, the last subframe can be padded with zeros.

Each subframe $f = 1, \dots, F$ is compressed, producing a shorter binary sequence, which is then channel encoded, obtaining $L_b L_{\text{enc}}$ samples $s_{i,m}^k$, with $i = (f-1)L_b L_{\text{enc}} + 1, \dots, f L_b L_{\text{enc}}$, for all transmit antennas $m = 1, \dots, N^T$. Recall that the bandwidth expansion factor L_b specifies the number of channel uses per sensed sample that are available on the communication channel. The choice of the subframe size L_{enc} generally entails a trade-off between latency due to subframe processing and communications, on the one hand, and the effectiveness of block encoding, on the other hand, which is known to improve as the blocklength increases [36].

Each discrete-time symbol in the sequence $s_{i,m}^k$ is modulated using a waveform of bandwidth $1/T_c$, and transmitted over the wireless shared channel (1) to the receiver. Based on the received signal (3) in each subframe f , the receiver performs channel decoding and source decoding to recover an estimate $\hat{\mathbf{u}}_f^k$ of the original f th subframe. A specific implementation of source and channel encoders/decoder is described in Section VI.

B. Inference

While the proposed framework can be applied more broadly, in this work, we focus on supervised learning for a classification. In this subsection, we describe how inference is implemented, while the next subsection covers training.

The inference task amounts to a mapping between the sensed input \mathbf{u}^k from all devices $k = 1, \dots, K$ and a target variable \mathbf{v} . Given our focus on classification, we consider \mathbf{v} to be a $D_v \times 1$ one-hot vector representing the class index, where D_v is the number of classes. For the conventional system under study, we introduce an ANN that takes as input the per-subframe estimated sensed signals $\hat{\mathbf{u}}_f^k$ produced by the source decoder for all $k = 1, \dots, K$ to output an estimate $\hat{\mathbf{v}}_f$ of the target vector. The decisions produced for each subframe f are then combined to output a final decision $\hat{\mathbf{v}}$, e.g., via a majority rule. This approach is useful to enable the ANN to make a decision after the reception of each subframe. Other conventional solutions may involve the use of Recurrent Neural Networks (RNNs), and they are not further considered here.

Mathematically, the ANN defines a parametric mapping $\hat{\mathbf{v}}_f = G_\phi(\hat{\mathbf{u}}_f)$, where ϕ is the vector of parameters of the ANN, and the output $\hat{\mathbf{v}}_f \in [0, 1]^{D_v}$ provides an estimate of the probabilities of the D_v classes, which is obtained via a softmax output layer. The decision at the f subframe can be made by choosing the class with the largest probability in vector $\hat{\mathbf{v}}_f$. After each subframe

f is received, a more accurate decision \hat{v} may be obtained through a fixed aggregation function $\text{agg}(\hat{v}_1, \dots, \hat{v}_f)$ of the individual per-subframe decision. This function may only account for the most recent subframes to facilitate online decisions. An example is given by the majority rule that outputs the class selected over the majority of aggregated subframes.

C. Training

In order to optimize the model parameter vector ϕ of the ANN, we assume access to a training data set \mathcal{D} containing input-output pairs (\mathbf{u}, \mathbf{v}) , where \mathbf{u} collects all the samples sensed by the K devices and \mathbf{v} is the corresponding target one-hot variable. The sensed signal \mathbf{u} is formatted into F subframes \mathbf{u}_f , with $f = 1, \dots, F$, as explained above. The loss between the probability vector $\hat{\mathbf{v}}_f = G_\phi(\mathbf{u}_f)$ and the true target \mathbf{v} is measured by the cross-entropy loss [37]

$$L(\hat{\mathbf{v}}_f, \mathbf{v}) = - \sum_{j=1}^{D_v} v_j \log(\hat{v}_{j,f}), \quad (8)$$

where v_j and $\hat{v}_{j,f}$ are the j th entry of vector \mathbf{v} and $\hat{\mathbf{v}}_f$, respectively. We wish to select the model parameter vector ϕ based on the training example (\mathbf{u}, \mathbf{v}) by addressing the empirical risk minimization (ERM) problem

$$\min_{\phi} \sum_{(\mathbf{u}, \mathbf{v}) \in \mathcal{D}} \sum_{f=1}^F L(G_\phi(\mathbf{u}_f), \mathbf{v}). \quad (9)$$

Accordingly, the model parameters ϕ can be updated based on stochastic gradient descent (SGD)-based rule

$$\phi = \phi - \eta \sum_{(\mathbf{u}, \mathbf{v}) \in \mathcal{B}} \sum_{f=1}^F \nabla_{\phi} L(G_\phi(\mathbf{u}_f), \mathbf{v}), \quad (10)$$

where $\mathcal{B} \subseteq \mathcal{D}$ is a mini-batch of data points and η is the learning rate.

IV. NEUROCOMM

In this section, we present the proposed end-to-end neuromorphic architecture for remote inference over wireless channels, which we refer to as NeuroComm. As introduced in Sec. I and detailed in Fig. 3, in NeuroComm, neuromorphic transmission and processing blocks are integrated with neuromorphic sensing. We proceed by describing neuromorphic encoding, decoding, and processing. Throughout this section, we consider the encoding and decoding SNNs to be fixed, i.e., to have fixed synaptic weights. In the next section, we will introduce the proposed hypernetwork-based architecture that enables optimization of encoding and decoding SNNs, as well as adaptation of the decoding SNN to current channel conditions.

A. Neuromorphic Encoding

With neuromorphic encoding, at each device k , the D_u samples \mathbf{u}_l^k produced by the sensor are encoded in an online, streaming, fashion by an SNN, across time instants $l = 1, 2, \dots, L$. The samples output by each encoding SNN are fed to the impulse radio transmitting module. The encoding SNNs hence implement *joint source-channel coding*, which is unlike the conventional solution reviewed in the previous section.

We investigate two distinct ways to implement neuromorphic encoding: one based on conventional *time hopping* (TH) for impulse radio (see, e.g., [19]), and the other on an optimized solution that we refer to as *learned time hopping* (LTH) modulation. As illustrated in Fig. 5, the distinction between the two schemes is in the way in which bandwidth expansion is handled. In TH, the k th encoding SNN produces a single sample x_l^k for each sensed sample u_l^k . Bandwidth expansion is then carried out via TH in order to produce the samples s_i^k to be transmitted on the wireless channel (3). Specifically, as detailed below, TH outputs L_b samples s_i^k , with $i \in \mathcal{I}_l = \{(l-1)L_b + 1, \dots, lL_b\}$, for each sample by randomizing the timing of any spike in x_l^k . In contrast, in LTH, the encoding SNNs directly produce the L_b samples s_i^k , with $i \in \mathcal{I}_l$, given each sensed sample u_l^k .

For both TH and LTH, the encoding SNN at each device k , with $k = 1, \dots, K$, defines a causal mapping between the sensed time sequence \mathbf{u}_l^k , and the encoded time sequence. The encoded time sequence – x_l^k for TH and s_i^k for LTH – consists of $D_x = N^T$ spiking signals. Each of the D_x signals is mapped to the output of one of the N^T transmit antennas. We note that one could accommodate a larger number of antennas with $N^T > D_x$ by applying space-time codes for pulse-position modulation [34], but we do not further pursue this direction in this work.

1) *TH Modulation*: As illustrated in Fig. 5-(a), with standard TH, the encoding SNN produces an $N^T \times 1$ vector of samples $x_{l,m}^k$ for $m = 1, \dots, N^T$ given each sensed sample u_l^k . Each sample $x_{l,m}^k$ is then expanded into L_b samples $s_{i,m}^k$, for $i \in \mathcal{I}_l$, by randomly adding a discrete-time offset $c_{l,m}^k \in \{1, \dots, L_b\}$ to the timing of each spike. For example, the spike $x_{1,m}^k = 1$ in Fig. 5-(a) hops to position $c_{1,m}^k = 2$ in the first interval of sensing time. The sequence of random hop times $c_{l,m}^k \in \{1, \dots, L_b\}$ is not assumed to be known to the receiver.

2) *LTH Modulation*: TH modulation is motivated by an attempt to avoid interference between waveforms transmitted by different antennas and devices by randomizing the position of a spike within a given interval of L_b samples. However, in NeuroComm the goal is not to recover the individual transmitted signals, but rather to carry out an inference task. Therefore, TH may not




Fig. 5. (a) *Time hopping (TH) modulation*: Each output sample $x_{l,m}^k$ produced by the encoding SNN for the m th antenna at device k is expanded to $L_b = 5$ samples via TH. Accordingly, a spike, i.e., a sample $x_{l,m}^k = 1$, is encoded as a single spike with a random time within the interval of L_b samples corresponding to the l th sensing time step. (b) *Learned time hopping (LTH) modulation*: The encoding SNN directly outputs L_b samples $s_{i,m}^k$ with $i \in \mathcal{I}_l = \{(l-1)L_b + 1, \dots, lL_b\}$, for each sensed sample u_l^k .

be the most effective solution. Based on this observation, we introduce LTH, whereby we let the SNN utilize the available L_b channel uses per sensed sample in a manner that is optimized (as we will study in the next section) for the given end-to-end task.

To this end, we expand the sensed signal u_l^k for each time instant l by adding $L_b - 1$ zero samples before feeding the sequence of samples to the encoding SNN. As a result, the encoding SNN directly produces L_b output samples $s_{i,m}^k$, with $i \in \mathcal{I}_l$, in each interval of sensing time l . These samples are directly transmitted on the channel (3) as discussed next.

B. Impulse Radio Transmission

The samples $s_{i,m}^k$ produced by the encoding SNN are modulated onto antenna m by device k using pulse modulation. Accordingly, each antenna outputs a stream of pulses as illustrated




Fig. 6. The impulse radio block takes as input the spiking signals produced by the encoding SNN, and possibly by TH modulation, and it produces the modulated signals $s_m^k(t)$, for $m = 1, \dots, N^T$.

in Fig. 6, and, unlike the conventional solution described in the previous section, information in NeuroComm is directly encoded in the timing of the spikes. Importantly, if the sequence $s_{i,m}^k$ is sparse, so is the signal transmitted by antenna m at device k . This way, NeuroComm naturally controls the transmitted energy as a function of the activity in the sensed signals. If the sensors produce no spikes, the encoding SNN is also idle, implying that minimal computing energy and no transmission energy are consumed. In this regard, we also observe that, with TH, at most a single pulse is produced in each period of $L_b T_c$ seconds, while with LTH one may have multiple pulses, up to L_b , in each such period.

C. Neuromorphic Decoding

The modulated signals from the K devices are transmitted over the channel, and the discrete-time received signals (3) are input to the decoding SNN, which operates at the same rate as the neuromorphic sensors. To feed the decoding SNN, we collect the received signals $\{y_{i,n}\}$ corresponding to the sensed time instant l , i.e., samples $y_{i,n}$ with $i \in \mathcal{I}_l$, into a $L_b \times 1$ vector $\mathbf{y}_{l,n}$. Since this vector is generally complex, we define the $2L_b \times 1$ vector

$$\bar{\mathbf{y}}_{l,n} = [\Re(\mathbf{y}_{l,n})^T, \Im(\mathbf{y}_{l,n})^T]^T, \quad (11)$$

with $\Re(\cdot)$ and $\Im(\cdot)$ being the element-wise real and imaginary parts of the input vector, respectively. The decoding SNN implements a causal mapping between the sequence of $2L_b N^R \times 1$ vectors $\bar{\mathbf{y}}_l = [\bar{\mathbf{y}}_{l,1}^T, \dots, \bar{\mathbf{y}}_{l,N^R}^T]^T$ and spiking signals $\bar{\mathbf{v}}_l \in \{0, 1\}^{D_v}$ across time instants $l = 1, \dots, L$.

Finally, the D_v spiking signals in vector sequence \bar{v}_l are used as input to a block that makes the final inference decision.

For a classification task with D_v classes, at each discrete time l , a decision can be made by decoding the output signal $[\bar{v}_1, \dots, \bar{v}_l]$ produced thus far by the decoding SNN to a class index using standard methods, such as rate decoding (see, e.g., [14]). With rate decoding, the D_v signals $[\bar{v}_1, \dots, \bar{v}_l]$ are first converted into a $D_v \times 1$ vector by counting the spikes in each of the D_v spiking signals, and then the $D_v \times 1$ vector of rates is passed through a softmax function. This can be written as

$$\hat{v}_l = \text{softmax}(\bar{\mathbf{V}}_l \times \mathbf{1}), \quad (12)$$

where $\bar{\mathbf{V}}_l$ is the $D_v \times l$ matrix $\bar{\mathbf{V}}_l = [\bar{v}_1, \dots, \bar{v}_l]$; $\mathbf{1}$ is the all-ones $l \times 1$ vector; and the softmax function outputs a $D_v \times 1$ vector with the j th entry given by $\text{softmax}_j(\mathbf{x}) = e^{x_j} / (\sum_{j=1}^{D_v} e^{x_j})$ for a $D_v \times 1$ input vector $\mathbf{x} = [x_1, \dots, x_{D_v}]$. Then, the class corresponding to the largest number of spikes across L time steps, that is, to the largest entry in vector \hat{v} , is selected.

D. Spiking Neural Network Model

In this work, each encoding/decoding SNN has a layered structure with fully connected layers. We describe here the general architecture using a notation that does not specify the identity of the SNN, with the understand that superscripts can be added to distinguish encoding and decoding SNNs. For instance, we will write as N_L the number of layers with N_L^k being the number of layer for the k th encoding SNN and N_L^R being the number of layer for the decoding SNN.

Each layer ℓ , with $\ell = 1, \dots, N_L$, consists of N_ℓ spiking neurons, and we denote by \mathbf{W}_ℓ the $N_{\ell+1} \times N_\ell$ weight matrix describing the connections between neurons in layer ℓ and layer $\ell + 1$. Neurons follow the standard spike response model (SRM) [38]. Each spiking neuron k outputs a binary signal $b_{k,l} \in \{0, 1\}$, with “1” representing the firing of the spike and “0” an idle neuron at each time step $l = 1, \dots, L$. Each neuron k in layer $\ell + 1$ receives inputs from the set \mathcal{N}_ℓ of neurons in layer ℓ .

Following the SRM, each neuron k maintains an internal analog state variable $o_{k,l}$, known as the *membrane potential*, over time step l . Mathematically, the membrane potential $o_{k,l}$ is defined by the sum of filtered contributions from incoming spikes and from the neuron k 's own past outputs. Mathematically, we write the evolution of the membrane potential as

$$o_{k,l} = \sum_{j \in \mathcal{N}_\ell} w_{k,j,\ell} \cdot (\alpha_l * b_{j,l}) + \beta_l * b_{k,l}, \quad (13)$$

where $w_{k,j,\ell}$ is the element (k,j) of matrix \mathbf{W}_ℓ , which corresponds to the synaptic weight between neuron $j \in \mathcal{N}_\ell$ and neuron k in layer $\ell + 1$; α_l represents the synaptic response to a spike from the presynaptic neurons $j \in \mathcal{N}_\ell$ to a postsynaptic neuron k ; β_l describes the synaptic response to the spike emitted by the neuron itself; and “ $*$ ” is the convolution operator. Typical choices for synaptic spike responses include the first-order feedback filter $\beta_l = \exp(-l/\tau_{\text{ref}})$, and the second-order synaptic filter $\alpha_l = \exp(-l/\tau_{\text{mem}}) - \exp(-l/\tau_{\text{syn}})$, for $l = 1, 2, \dots$, with finite positive constants τ_{ref} , τ_{mem} and τ_{syn} .

Neuron k outputs a spike at time step l when its membrane potential $o_{k,l}$ passes some fixed threshold ϑ , i.e.,

$$b_{k,l} = \Theta(o_{k,l} - \vartheta), \quad (14)$$

where $\Theta(\cdot)$ is the Heaviside step function.

V. ENABLING ADAPTATION VIA HYPER-NEUROCOMM

Deploying NeuroComm requires the specification of the synaptic weights $\mathbf{W}^1, \dots, \mathbf{W}^K$ of the K encoding SNNs, as well as of the weights \mathbf{W}^R of the decoding SNN. In this section, we introduce an end-to-end design strategy for all the synaptic weights. Design is carried out offline based on the availability of channel models or channel data. Importantly, the proposed approach does not require retraining for each channel condition. Rather, optimization of the synaptic weights is done only once, but the system is still able to adjust to changing channel conditions. Specifically, in the proposed architecture, the encoding SNNs’ weights are fixed irrespective of the current channel conditions, which are assumed to be unknown to the transmitters. In contrast, we allow for the weights \mathbf{W}^R of the decoding SNN to adapt to the current channel condition by leveraging the received pilots (see Section II).

To this end, we propose a hybrid receiver architecture, termed Hyper-NeuroComm, which includes a conventional hypernetwork. As illustrated in Fig. 7, the hypernetwork takes as input the pilots received in the current frame, and outputs information controlling the weights of the decoding SNN. During offline training, the transmitters’ weights, as well as the receiver’s adaptation procedure via the hypernetwork, are jointly optimized. In this section, we first describe the Hyper-NeuroComm architecture, and then define the end-to-end training problem for the design of the encoding SNNs, decoding SNN, and hypernetwork.

A. Hyper-NeuroComm

Hypernetworks are neural networks whose outputs determine the weights of another neural network [39]. The target network in our setting is the decoding SNN. Accordingly, the ideal goal of the hypernetwork is to implement a parametric function

$$\mathbf{W}^R = h_{\mathbf{W}^H}^{\text{id}}(\bar{\mathbf{y}}_p) \quad (15)$$

that maps the received pilots signals $\bar{\mathbf{y}}_p$ in (6) to the weights \mathbf{W}^R of the decoding SNN. The mapping (15) is parametrized by the weights \mathbf{W}^H of the hypernetwork. Via (15), the weights \mathbf{W}^R of the decoding SNN can be adapted to the current channel conditions thanks to the received pilots $\bar{\mathbf{y}}_p$. More practically, following [32], we decrease the dimension of the output (15) of the hypernetwork by producing a common scaling factor for each neuron, rather than a different weight for each synapse as in (15). The overall architecture of the receiver is shown in Fig. 7, and is detailed next.

We start by collecting all the received pilot signals into a vector $\mathbf{y}_p = [y_{-L_p+1,1}^1, \dots, y_{i,n}^k, \dots, y_{0,N^R}^K]^T \in \mathbb{C}^{L_p K N^R \times 1}$, and take the $2L_p K N^R \times 1$ vector $\bar{\mathbf{y}}_p = [\Re(\mathbf{y}_p)^T, \Im(\mathbf{y}_p)^T]^T$ as input to the hypernetwork. The output of the hypernetwork consists of N_L^R vectors $\{\mathbf{w}_1^R, \dots, \mathbf{w}_{N_L^R}^R\}$, with each $N_\ell^R \times 1$ vector \mathbf{w}_ℓ^R defining the scaling of the weights applied by the N_ℓ^R neurons in the ℓ th layer. Specifically, following [32], the weight matrix \mathbf{W}_ℓ^R for layer ℓ in the decoding SNN is set as

$$\mathbf{W}_\ell^R = \tilde{\mathbf{W}}_\ell^R \cdot \text{diag}\{\mathbf{w}_\ell^R\}, \quad (16)$$

for $\ell = 1, \dots, N_L^R$, where $\text{diag}\{\mathbf{w}_\ell^R\}$ is the $N_\ell^R \times N_\ell^R$ diagonal matrix with main diagonal given by the elements of $N_\ell^R \times 1$ vector \mathbf{w}_ℓ^R ; and the $N_{\ell+1}^R \times N_\ell^R$ matrices $\tilde{\mathbf{W}}_\ell^R$ for all ℓ are subject to optimization during training and they are fixed at run time, as the encoding SNNs' weights $\mathbf{W}^1, \dots, \mathbf{W}^K$ (see next subsection). We write $\tilde{\mathbf{W}}^R$ for the collection of all matrices $\{\tilde{\mathbf{W}}_\ell^R\}_{\ell=1}^{N_L^R}$, and the overall mapping implemented by the hypernetwork as

$$\tilde{\mathbf{W}}^R = h_{\mathbf{W}^H}(\bar{\mathbf{y}}_p), \quad (17)$$

where \mathbf{W}^H are the weights of the hypernetwork.

B. Training Problem

In this subsection, we address the problem of optimizing the model parameters $\mathbf{W}^1, \dots, \mathbf{W}^K$ of the encoding SNNs, along with the weights $\tilde{\mathbf{W}}^R$ in (16) of the decoding SNN, the weights \mathbf{W}^H




Fig. 7. The Hyper-NeuroComm architecture, enabling the adaptation of the weights of the decoding SNN.

of the hypernetwork in (17), and the pilot sequences s_p^1, \dots, s_p^K . We write as $\theta = [\mathbf{W}^1, \dots, \mathbf{W}^K, \tilde{\mathbf{W}}^R, \mathbf{W}^H, s_p^1, \dots, s_p^K]$ the set of all parameters under optimization. As in Section III-C, we assume the availability of a training data set \mathcal{D} given by pairs of inputs $\mathbf{u} = (\mathbf{u}^1, \dots, \mathbf{u}^K)$ produced by all sensors across all L time instants, and the corresponding target one-hot vector \mathbf{v} . As described in Section IV-C, the decoding SNN outputs spiking signals that are converted into a $D_v \times 1$ probability vector $\hat{\mathbf{v}}_l$ via (12). The output $\hat{\mathbf{v}}_l$ depends on the input \mathbf{u} through the encoding SNNs' weights $\mathbf{W}^1, \dots, \mathbf{W}^K$, the decoding SNN's weights \mathbf{W}^R , the channels $\mathbf{h} = \{ \{ \{ \mathbf{h}_{n,m}^k \}_{n=1}^{N^R} \}_{m=1}^{N^R} \}_{k=1}^K$, as well as the additive channel noise in (3).

Using the cross-entropy loss (8), the training objective for Hyper-NeuroComm is defined as

$$\min_{\theta} \sum_{(\mathbf{u}, \mathbf{v}) \in \mathcal{D}} \mathbb{E}[L(\hat{\mathbf{v}}_L, \mathbf{v})], \quad (18)$$

where $\mathbb{E}[\cdot]$ denotes the expectation operator over channels \mathbf{h} and noise, and we recall that the decoding SNN's weights \mathbf{W}^R depend on the hypernetwork's weights \mathbf{W}^H as in (16). In practice, the expectation in (18) is approximated by sampling channel realizations \mathbf{h} and channel noise.

Problem (18) is addressed via stochastic gradient descent in a manner similar to (10). In order to compute the gradient, we follow the surrogate gradient approach [40] by replacing the Heaviside step function (14) with a differentiable surrogate function, namely a sigmoid $\sigma(x) = (1 + e^{-x})^{-1}$.

VI. EXPERIMENTS

In this section, we provide experimental results with the aim of comparing the performance of the proposed end-to-end neuromorphic solution – Hyper-NeuroComm – with more conventional methods. As we detail next, we study a remote inference setting in which devices are equipped with DVS cameras observing parts of a screen displaying a handwritten digit. The goal is to ensure that the receiver be able to quickly classify the displayed image.



Fig. 8. An illustration of a system with two devices ($K = 2$), each equipped with a DVS camera that observes a fraction μ of an image displayed on a screen. In the depicted example, the fraction μ is smaller than 0.5, and hence the two devices cannot collectively observe the whole image.

A. Experimental Setting

To set up the scenario, we consider the MNIST-DVS dataset [41], which contains labelled 26×26 spiking signals of duration $L = 80$ samples. Each data point contains $26 \cdot 26 = 676$ spiking signals, which are recorded from a DVS camera that is shown moving handwritten digits “1” and “7” on a screen. The data set contains 900 examples from each class for training, and 100 examples for testing. The neuromorphic sensor at each device records a fraction $\mu < 1$ of the image. Accordingly, each device has an input signal consisting of $D_u = \lceil \mu \cdot 676 \rceil$ spiking signals. Specifically, for the case of $K = 2$ sensors, we assume, as shown in Fig. 8, that the first device observes the fraction of top μ rows of the image, while the second device measures the bottom μ -fraction of the rows. Note that when $\mu > 1/2$ there is an overlap of size $2\mu - 1$ between the subimages observed by the two sensors.

Each sensor has $N^T = 10$ antennas, and the decoder has $N^R = 20$ antennas. The encoding SNNs has 128 hidden neurons, and $D_x = N^T = 10$ output neurons. The decoder has $D_y = 20$

input neurons, 128 hidden neurons, and $D_v = 2$ output neurons, each corresponding to one of the two classes. The hypernetwork ANN is composed of one hidden layer with 1024 neurons. All layers are fully connected in the SNNs and in the hypernetwork. The pilots are initialized as white Gaussian noise and then subject to optimization during training by tackling problem (18). The number of pilots L_p is set to 64.

The transmitter uses Gaussian monopulses, and the receiver performs classification at each discrete time instant. The number of paths N_P is set to 5, all paths have the same average power $1/N_P$, and they are characterized by Rayleigh fading. The multipath delays are set to $0, T_c, 2T_c, 3T_c$ and $4T_c$. Unless otherwise stated, we consider the per-frame power constraint (4), and the average signal-to-noise ratio (SNR) for each device k is defined as the average symbol power over noise power $(E_{\text{fr}}^k / (D_x L L_b)) / N_0$. When considering per-symbol energy constraint (4), the SNR is defined as E_s^k / N_0 . The average SNR is set to 10 dB. We train all the schemes based on the MNIST-DVS dataset with 1000 channel and noise realizations for each sample, while the test results are averaged over 300 channel realizations.

B. Benchmarks

Given that this is the first work to propose the idea of neuromorphic computing for remote inference, there are no direct benchmarks involving neuromorphic technologies. Therefore, in order to benchmark the proposed approach, we rely on state-of-the-art conventional frame-based strategies that adopt traditional multi-access technologies and separate source and channel coding. We also consider several simplified variants of the proposed neuromorphic architecture.

For the proposed Hyper-NeuroComm, we implement both standard TH and LTH (see Sec. IV). Furthermore, we consider the following benchmarks.

- *Frame-based digital transmission:* Following the description in Section III, we implement frame-based digital transmission by following [24]. The transmitter collects $L_{\text{enc}} = 20$ time samples from the sensor to create a subframe. The subframe is encoded by source and channel encoding separately. Source encoding is performed using a Vector-Quantization Variational Autoencoder (VQ-VAE) [42], followed by LDPC channel encoding with code rate set to $2/3$. The encoded data is transmitted subframe by subframe using BPSK to the receiver. Each subframe is decoded using the belief propagation algorithm, decompressed by the VQ-VAE, and then classified by an ANN, which consists of a single layer of 512 hidden neurons. Frame-based digital transmission is implemented with ALOHA to manage

interference between the transmitted signals from these two devices. Accordingly, each device transmits a subframe with probability 0.5, and subframe transmission is successful only when there are no concurrent transmissions. The aggregation function adopts a majority rule over the received subframes. The DNN is trained using SGD-based optimization based on the examples (\mathbf{u}, \mathbf{v}) as per (10).

- *NeuroComm with joint learning*: To set a lower bound on the performance of the proposed Hyper-NeuroComm architecture, we consider an alternative design method whereby the decoding SNN model parameters are not adapted to the channel realization using pilots. This solution is derived by removing the hypernetwork, and addressing problem (18) directly over encoding and decoding SNN's weights.
- *NeuroComm with ideal per-channel learning*: To provide a performance upper bound, we also consider an ideal situation in which for each given channel realization we have abundant data to train the model parameter θ to convergence. This is done by tackling problem (18) via SGD by drawing a large number of channel and noise samples.

C. Results: Single Device

We first consider a baseline scenario in which there is only one device observing the entire image \mathbf{u} , i.e., we have $K = 1$ and $\mu = 1$. We recall that reference [24] also considered the scenario with $K = 1$ and $\mu = 1$, but only under a simple additive Gaussian noise channel with no fading. Therefore, the performance analysis reported here, which accounts for the impact of fading and the corresponding novel role played by the pilots and by the hypernetwork, is novel. We first focus on LTH in Fig. 9, which generally outperform TH, and then compare LTH and TH in Fig. 10.

In Fig. 9, we evaluate the evolution of the test accuracy during the observation of the signals from the sensors. The horizontal axis reports the discrete time l elapsed since the start of the observation of the image by the sensors. The test accuracy of all schemes is seen to improve over time, as they can make decision based on more samples. However, the frame-based approach enables an improvement of the accuracy level only after each subframe is received, decoded, and processed by the DNN. NeuroComm with joint learning performs slightly better than frame-based transmission, offering a gradual increase in test accuracy. However, the gain is limited by the lack of adaptation to changing conditions.




Fig. 9. Test accuracy versus discrete time l during transmission of an image with LTH ($K = 1$, $\mu = 1$, $L_b = 6$, $SNR = 10$ dB).

Hyper-NeuroComm is seen to outperform NeuroComm with joint learning, achieving a test accuracy close to the ideal case of per-channel learning. This demonstrates the advantages of the proposed hypernetwork-based architecture in transferring useful knowledge across fading realizations, while also being able to adapt to current channel conditions.

Fig. 10 shows the test accuracy as a function of the bandwidth expansion factor L_b . We recall that a larger value of L_b provides the encoding SNN with a proportionally larger interval of time over which to encode the sensed input. The test accuracy is observed to generally increase with the bandwidth expansion L_b , owing to the decreased interference between the transmitted streams from the N^T antennas. Furthermore, LTH is seen to be effective in making a more efficient use of the increased encoding period. Rather than only managing interference between streams, LTH can in fact effectively use the larger number of encoded symbols per sensed sample to enhance the efficiency of joint source-channel encoding from sensed to transmitted signals.

Finally, Fig. 11 demonstrates the cumulative transmitted energy consumption as a function of the time step l for frame-based digital transmission and Hyper-NeuroComm. Unlike the previous figures, we adopt the per-symbol energy constraint (5). The images are presented starting from time step 10. While frame-based digital transmission scheme consumes constant




Fig. 10. Test accuracy versus L_b for NeuroComm ($K = 1$, $\mu = 1$, $SNR = 10$ dB).




Fig. 11. Cumulative transmitted energy consumption versus time step l under a per-symbol energy constraint ($K = 1$, $\mu = 1$, $L_b = 5$, $SNR = 10$ dB).

power, NeuroComm only consumes energy when images are presented to the DVS cameras.




Fig. 12. Test accuracy versus fraction μ of observed image during transmission of an image with LTH ($K = 2$, $SNR = 10$ dB).

Furthermore, the cumulative energy consumption is lower than for frame-based transmission, due to the sparsity of the spiking signals output by the sensors and by the encoding SNNs. Finally, the LTH scheme consumes more energy than TH, since the LTH may produce multiple pulses during each time step.

D. Results: Multiple Devices

We now move on to a setting with $K = 2$ devices, each observing a fraction μ of the screen as described in Fig. 8. In Fig. 12, we shows the test accuracy as a function of the fraction μ of the sensed signals under LTH modulation. Test accuracy is seen to improve as sensors record a larger fraction of an image. However, only Hyper-NeuroComm is able to fully benefit from the more informative sensed signals, achieving a performance close to ideal per-channel learning.

Fig. 13 plots the test accuracy versus the bandwidth expansion factor L_b for Hyper-NeuroComm with $K = 1$ and $K = 2$ devices under LTH and TH. An interesting question arising from this experiment is: Under which conditions does having two sensors help? The figure shows that, if the bandwidth expansion is too low, e.g., if $L_b < 4$ for LTH, it is preferable to have a single device active, despite the fact that the device only observes a fraction μ of the entire image.




Fig. 13. Test accuracy versus L_b for Hyper-NeuroComm ($\mu = 0.75$, $SNR = 10$ dB).

Furthermore, it can be noted that, while in the single-device case Hyper-NeuroComm with LTH achieves its best performance with a smaller value of L_b , e.g., $L_b = 9$, a larger L_b is called for in order to mitigate interference from multiple devices.

VII. CONCLUSION

This paper has introduced NeuroComm, a novel neuromorphic end-to-end solution for remote wireless inference in applications involving always-on online monitoring. The proposed system integrates neuromorphic sensing and neuromorphic computing with impulse radio transmission, enabling the adaptation of energy consumption to the dynamics of the monitored environment while ensuring low-latency inference. At a system level, the innovations of this paper also include the use of a hypernetwork at the receiver side to optimize adaptation of the receiver to the current channel conditions. Extensive experimental results have demonstrated the benefits of neuromorphic sensing and computing for the implementation of wireless cognition in terms of time-to-accuracy and energy consumption. It is the hope of the authors that this work will pave the way for further study of the synergy between neuromorphic sensing and processing technologies with wireless communications. A particularly interesting avenue for research would be the development of a testbed incorporating software-defined radio that implement IR, as well

as neuromorphic hardware, possibly accessed using cloud services such as Intel’s Lava [43]. This would make it possible to assess the relative contributions of processing and transmit-receive chains on the overall energy consumption. Furthermore, one could consider other types of inference tasks such as tracking [16]. Finally, the IR interface could be leveraged for conventional data transmission as well as a probing signal for sensing applications [44].

REFERENCES

- [1] P. Popovski, O. Simeone, F. Boccardi, D. Gündüz, and O. Sahin, “Semantic-effectiveness filtering and control for post-5G wireless connectivity,” *Journal of the Indian Institute of Science*, vol. 100, no. 2, pp. 435–443, 2020.
- [2] E. C. Strinati and S. Barbarossa, “6G networks: beyond shannon towards semantic and goal-oriented communications,” *Computer Networks*, vol. 190, p. 107930, 2021.
- [3] M. Kountouris and N. Pappas, “Semantics-empowered communication for networked intelligent systems,” *IEEE Communications Magazine*, vol. 59, no. 6, pp. 96–102, 2021.
- [4] M. Kalfa, M. Gok, A. Atalik, B. Tegin, T. M. Duman, and O. Arıkan, “Towards goal-oriented semantic signal processing: applications and future challenges,” *Digital Signal Processing*, vol. 119, p. 103134, 2021.
- [5] Z. Qin, X. Tao, J. Lu, and G. Y. Li, “Semantic communications: principles and challenges,” *arXiv preprint arXiv:2201.01389*, 2021.
- [6] F. Conti, P. D. Schiavone, and L. Benini, “XNOR neural engine: a hardware accelerator IP for 21.6-fJ/op binary neural network inference,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 37, no. 11, pp. 2940–2951, 2018.
- [7] A. Mehonic and A. J. Kenyon, “Brain-inspired computing needs a master plan,” *Nature*, vol. 604, no. 7905, pp. 255–260, 2022.
- [8] M. Davies, A. Wild, G. Orchard, Y. Sandamirskaya, G. A. F. Guerra, P. Joshi, P. Plank, and S. R. Risbud, “Advancing neuromorphic computing with Loihi: a survey of results and outlook,” *Proceedings of the IEEE*, vol. 109, no. 5, pp. 911–934, 2021.
- [9] “inivation ag,” [Online], <https://inivation.com/>.
- [10] Prophesee, “Event-based vision applications,” [Online], <https://www.prophesee.ai/>.
- [11] H. Rebecq, R. Ranftl, V. Koltun, and D. Scaramuzza, “Events-to-video: bringing modern computer vision to event cameras,” in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3852–3861, 2019.
- [12] W. W. Lee, Y. J. Tan, H. Yao, S. Li, H. H. See, M. Hon, K. A. Ng, B. Xiong, J. S. Ho, and B. C. Tee, “A neuro-inspired artificial peripheral nervous system for scalable electronic skins,” *Science Robotics*, vol. 4, no. 32, 2019.
- [13] S. Ward-Foxton, “Neuromorphic developers partner to integrate sensor, processor,” [Online], <https://www.eetimes.com/neuromorphic-developers-partner-to-integrate-sensor-processor/>.

- [14] H. Jang, O. Simeone, B. Gardner, and A. Gruning, "An introduction to probabilistic spiking neural networks: probabilistic models, learning rules, and applications," *IEEE Signal Processing Magazine*, vol. 36, no. 6, pp. 64–77, 2019.
- [15] B. Rajendran, A. Sebastian, M. Schmuken, N. Srinivasa, and E. Eleftheriou, "Low-power neuromorphic hardware for signal processing applications: a review of architectural and system-level design approaches," *IEEE Signal Processing Magazine*, vol. 36, no. 6, pp. 97–110, 2019.
- [16] Y. Hu, H. Liu, M. Pfeiffer, and T. Delbruck, "DVS benchmark datasets for object tracking, action recognition, and object recognition," *Frontiers in neuroscience*, vol. 10, p. 405, 2016.
- [17] P. Lichtsteiner, C. Posch, and T. Delbruck, "A 128 x 128 120db 30mw asynchronous vision sensor that responds to relative intensity change," in *2006 IEEE International Solid State Circuits Conference-Digest of Technical Papers*, pp. 2060–2069, 2006.
- [18] M. Davies, N. Srinivasa, T.-H. Lin, G. China, 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.
- [19] M. Z. Win and R. A. Scholtz, "Impulse radio: how it works," *IEEE Communications letters*, vol. 2, no. 2, pp. 36–38, 1998.
- [20] A. Shahshahani, M. Shahshahani, P. M. Ros, A. Bonanno, M. Crepaldi, M. Martina, D. Demarchi, and G. Masera, "An all-digital spike-based ultra-low-power IR-UWB dynamic average threshold crossing scheme for muscle force wireless transmission," in *2015 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pp. 1479–1484, 2015.
- [21] M. Z. Win and R. A. Scholtz, "Ultra-wide bandwidth time-hopping spread-spectrum impulse radio for wireless multiple-access communications," *IEEE Transactions on communications*, vol. 48, no. 4, pp. 679–689, 2000.
- [22] I. F. Akyildiz, C. Han, Z. Hu, S. Nie, and J. M. Jornet, "Terahertz band communication: an old problem revisited and research directions for the next decade," *IEEE Transactions on Communications*, 2022.
- [23] P. Yang, Y. Xiao, M. Xiao, and S. Li, "6G wireless communications: vision and potential techniques," *IEEE Network*, vol. 33, no. 4, pp. 70–75, 2019.
- [24] N. Skatchkovsky, H. Jang, and O. Simeone, "End-to-end learning of neuromorphic wireless systems for low-power edge artificial intelligence," in *Proc. Asilomar Conference on Signals, Systems, and Computers*, pp. 166–173, 2020.
- [25] T. Borsos, M. Condoluci, M. Daoutis, P. Haga, and A. Veres, "Resilience analysis of distributed wireless spiking neural networks," in *Proc. IEEE Wireless Communications and Networking Conference (WCNC)*, pp. 2375–2380, 2022.
- [26] E. M. Bansbach, A. Bank, and L. Schmalen, "Spiking neural network decision feedback equalization," *arXiv preprint arXiv:2211.04756*, 2022.
- [27] E. Bourtsoulatze, D. B. Kurka, and D. Gunduz, "Deep joint source-channel coding for wireless image transmission," *IEEE Transactions on Cognitive Communications and Networking*, vol. 5, no. 3, pp. 567–579,

- 2019.
- [28] M. Yang, C. Bian, and H.-S. Kim, “Deep joint source channel coding for wireless image transmission with OFDM,” in *Proc. IEEE International Conference on Communications (ICC)*, pp. 1–6, 2021.
 - [29] D. B. Kurka and D. Gündüz, “Bandwidth-agile image transmission with deep joint source-channel coding,” *IEEE Transactions on Wireless Communications*, vol. 20, no. 12, pp. 8081–8095, 2021.
 - [30] M. A. Alawad, M. Q. Hamdan, and K. A. Hamdi, “End-to-end deep Learning IRS-assisted communications systems,” in *Proc. IEEE Vehicular Technology Conference (VTC2021-Fall)*, pp. 1–6, 2021.
 - [31] M. Goutay, F. A. Aoudia, and J. Hoydis, “Deep hypernetwork-based MIMO detection,” in *Proc. IEEE International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, pp. 1–5, 2020.
 - [32] Y. Liu and O. Simeone, “HyperRNN: deep learning-aided downlink CSI acquisition via partial channel reciprocity for FDD massive MIMO,” in *Proc. IEEE International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, pp. 31–35, 2021.
 - [33] N. Skatchkovsky, O. Simeone, and H. Jang, “Learning to time-decode in spiking neural networks through the information bottleneck,” *arXiv preprint arXiv:2106.01177*, 2021.
 - [34] C. Han, A. O. Bicen, and I. F. Akyildiz, “Multi-wideband waveform design for distance-adaptive wireless communications in the terahertz band,” *IEEE Transactions on Signal Processing*, vol. 64, no. 4, pp. 910–922, 2015.
 - [35] I. M. Jacobs and J. Wozencraft, *Principles of communication engineering*. Waveland Press, 1965.
 - [36] M. Thomas and A. T. Joy, *Elements of information theory*. Wiley-Interscience, 2006.
 - [37] O. Simeone, *Machine learning for engineers*. Cambridge University Press, 2022.
 - [38] N. Skatchkovsky, H. Jang, and O. Simeone, “Spiking neural networks—part II: detecting spatio-temporal patterns,” *IEEE Communications Letters*, vol. 25, no. 6, pp. 1741–1745, 2021.
 - [39] D. Ha, A. Dai, and Q. V. Le, “Hypernetworks,” *arXiv preprint arXiv:1609.09106*, 2016.
 - [40] E. O. Neftci, H. Mostafa, and F. Zenke, “Surrogate gradient learning in spiking neural networks: bringing the power of gradient-based optimization to spiking neural networks,” *IEEE Signal Processing Magazine*, vol. 36, no. 6, pp. 51–63, 2019.
 - [41] T. Serrano-Gotarredona and B. Linares-Barranco, “Poker-DVS and MNIST-DVS: their history, how they were made, and other details,” *Frontiers in neuroscience*, vol. 9, p. 481, 2015.
 - [42] A. Van Den Oord, O. Vinyals *et al.*, “Neural discrete representation learning,” *Advances in neural information processing systems*, vol. 30, 2017.
 - [43] I. Corporation, “Lava software framework,” <https://lava-nc.org/>, accessed: 2022-11-08.
 - [44] J. Chen, N. Skatchkovsky, and O. Simeone, “Neuromorphic integrated sensing and communications,” *IEEE Wireless Communications Letters*, pp. 1–1, 2022.