

# Density Encoding Enables Resource-Efficient Randomly Connected Neural Networks

Denis Kleyko<sup>1</sup>, Member, IEEE, Mansour Kheffache<sup>2</sup>, E. Paxon Frady, Urban Wiklund<sup>3</sup>, Member, IEEE, and Evgeny Osipov<sup>4</sup>, Associate Member, IEEE

**Abstract**—The deployment of machine learning algorithms on resource-constrained edge devices is an important challenge from both theoretical and applied points of view. In this brief, we focus on resource-efficient randomly connected neural networks known as random vector functional link (RVFL) networks since their simple design and extremely fast training time make them very attractive for solving many applied classification tasks. We propose to represent input features via the density-based encoding known in the area of stochastic computing and use the operations of binding and bundling from the area of hyperdimensional computing for obtaining the activations of the hidden neurons. Using a collection of 121 real-world data sets from the UCI machine learning repository, we empirically show that the proposed approach demonstrates higher average accuracy than the conventional RVFL. We also demonstrate that it is possible to represent the readout matrix using only integers in a limited range with minimal loss in the accuracy. In this case, the proposed approach operates only on small  $n$ -bits integers, which results in a computationally efficient architecture. Finally, through hardware field-programmable gate array (FPGA) implementations, we show that such an approach consumes approximately 11 times less energy than that of the conventional RVFL.

**Index Terms**—Density-based encoding, hyperdimensional computing, random vector functional link (RVFL) networks.

## I. INTRODUCTION

An ability to provide insights and predictive analytics in real time is the greatest demand from businesses and industries to data-driven technologies. The vector of the current development targets enabling machine learning applications on connected devices (edge computing), such as smartphones, robots, and vehicles. The benefits of computing at the edge are tremendous: higher reliability of solutions due to the decoupling from the network connectivity and

Manuscript received February 10, 2020; revised July 2, 2020; accepted August 8, 2020. Date of publication August 24, 2020; date of current version August 4, 2021. This work was supported in part by the Swedish Research Council under Grant 2015-04677. The work of Denis Kleyko was supported in part by the European Union's Horizon 2020 Research and Innovation Programme under the Marie Skłodowska-Curie Individual Fellowship Grant Agreement 839179 and in part by the DARPA's VIP Program under Super-HD Project. (*Corresponding author: Denis Kleyko.*)

Denis Kleyko is with the Redwood Center for Theoretical Neuroscience, University of California at Berkeley, Berkeley, CA 94720 USA, and also with the Intelligent Systems Lab, Research Institutes of Sweden, 164 40 Kista, Sweden (e-mail: denis.kleyko@ri.se).

Mansour Kheffache is with Netlight Consulting AB, 111 53 Stockholm, Sweden (e-mail: mansour.kheffache@netlight.com).

E. Paxon Frady is with the Redwood Center for Theoretical Neuroscience, University of California at Berkeley, Berkeley, CA 94720 USA (e-mail: epaxon@berkeley.edu).

Urban Wiklund is with the Department of Radiation Sciences, Biomedical Engineering, Umeå University, 901 87 Umeå, Sweden (e-mail: urban.wiklund@umu.se).

Evgeny Osipov is with the Department of Computer Science Electrical and Space Engineering, Luleå University of Technology, 971 87 Luleå, Sweden (e-mail: evgeny.osipov@ltu.se).

This article has supplementary downloadable material available at <https://ieeexplore.ieee.org>, provided by the authors.

Color versions of one or more of the figures in this article are available online at <https://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNNLS.2020.3015971

bandwidth availability; very low latency; and higher security and privacy as sensitive data are processed locally on a device.

Randomly connected neural networks, such as the recently proposed class of advanced randomized learning techniques called stochastic configuration networks [1] and the well-known random vector functional link (RVFL) [2], have become an increasingly popular topic of modern theoretical and applied research. On the theoretical side, the main result is that RVFLs provide a universal approximation for continuous maps and functional approximations that converge in the Kullback–Leibler divergence when the target function is a probability density function [3]. When this is combined with the simplicity of RVFL's design and training process, it makes them a very attractive alternative for solving practical machine learning problems in edge computing.

The aim of this brief is to present an approach for an order of magnitude increase of the resource-efficiency (memory footprint, computational complexity, and energy consumption) of RVFLs operations. The proposed approach combines techniques from two fields of computer science: stochastic computing [4] and hyperdimensional computing [5]. The fundamental idea is in the realization of activations of the hidden layer with the computationally simple operations of hyperdimensional computing and the usage of the density-based encoding of the input features as in stochastic computing. Moreover, we enhance this approach with the integer-only readout matrix. This combination allows us to use integer arithmetics end-to-end. The novel contributions of this brief are given as follows.

- 1) A resource-efficient approach to RVFLs is proposed, which uses only integer operations.
- 2) The empirical evaluation on 121 real-world classification data sets demonstrates that the accuracy of the proposed approach is higher than that of the conventional RVFL.
- 3) Field-programmable gate array (FPGA) implementation of the proposed approach is an order of magnitude more energy-efficient and 2.5 times faster than the conventional RVFL.<sup>1</sup>

This brief is structured as follows. The background of methods used for the proposed approach is presented in Section II. The approach itself is described in Section III. The performance evaluation follows in Section IV. Section V covers related work. Section VI presents the concluding remarks.

## II. BACKGROUND AND METHODS

### A. Random Vector Functional Link

This section briefly describes the conventional RVFL. For a detailed survey of RVFLs, diligent readers are referred to [6]. Fig. 1 depicts the architecture of the conventional RVFL, which includes three layers of neurons. The input layer with  $K$  neurons represents the current values of input features denoted as  $\mathbf{x} \in [K \times 1]$ . The output layer ( $L$  neurons) produces the prediction of the network (denoted as  $\mathbf{y}$ ) during the operational phase. The layer in the middle

<sup>1</sup>For a network with 16 features, four classes, and 512 hidden neurons, which are the median values for the considered 121 data sets.




Fig. 1. Architecture of the conventional RVFL. In the presented example, the number of hidden neurons is set to  $N = 4$ .

is the hidden layer of the network, which performs a nonlinear transformation of input features. The hidden layer contains  $N$  neurons, and its state is denoted as  $\mathbf{h}$ .

In general, the connectivity of an RVFL is described by two matrices and a vector. A matrix  $\mathbf{W}^{\text{in}} \in [N \times K]$  describes connections between the input layer neurons and the hidden layer neurons. This matrix projects the given input features to the hidden layer. Each neuron in the hidden layer has a parameter called a bias. Biases of the hidden layer are stored in a vector and denoted as  $\mathbf{b} \in [N \times 1]$ . The other matrix of readout connections  $\mathbf{W}^{\text{out}} \in [L \times N]$  between the hidden and output layers transforms the current activations in the hidden layer stored in  $\mathbf{h}$  into the network's output  $\mathbf{y}$ .<sup>2</sup>

The main feature of the RVFL is that matrix  $\mathbf{W}^{\text{in}}$  and vector  $\mathbf{b}$  are randomly generated at the network initialization and stay fixed during the network's lifetime. There are no strict limitations for the generation of  $\mathbf{W}^{\text{in}}$  and  $\mathbf{b}$ . They are usually randomly drawn from either normal or uniform distributions. Here, both  $\mathbf{W}^{\text{in}}$  and  $\mathbf{b}$  are generated from a uniform distribution. Following [8], the range for  $\mathbf{W}^{\text{in}}$  is  $[-1, 1]$ , while the range for  $\mathbf{b}$  is  $[-0.1, 0.1]$ . Since  $\mathbf{W}^{\text{in}}$  and  $\mathbf{b}$  are fixed, the process of training RVFL is focused on learning the values of the readout matrix  $\mathbf{W}^{\text{out}}$ . The main advantage of training only  $\mathbf{W}^{\text{out}}$  is that the corresponding optimization problem is strictly convex; thus, the solution could be found in a single analytical step.

The activations of the network's hidden layer  $\mathbf{h}$  are described by the following equation:

$$\mathbf{h} = g(\mathbf{W}^{\text{in}}\mathbf{x} + \mathbf{b}) \quad (1)$$

where  $g(x)$  is a nonlinear activation function applied to each neuron. Here, the sigmoid function  $g(x) = (1/1 + e^{-x})$  is used. Thus, the activation function restricts the range of possible activation values in the hidden layer to the range  $[0, 1]$ .

The predictions issued by the output layer are calculated as

$$\mathbf{y} = \mathbf{W}^{\text{out}}\mathbf{h}. \quad (2)$$

With respect to the training of RVFLs, this brief focuses on classification tasks<sup>3</sup> considering only supervised-learning scenarios when the network is provided with the ground-truth label for each training example. The total size of the training data set is denoted as  $M$ . In this setting, the standard way of acquiring weights of the

<sup>2</sup>Strictly speaking, in the most general case, the readout matrix could also include connections between the input layer and the output layer. However, in the scope of this study, we only consider the case when the output layer predictions are obtained from the activations of the hidden layer. The interested readers are referred to work [7], which performed a comprehensive evaluation of different design choices for the RVFL.

<sup>3</sup>Though the proposed approach is also applicable to regression problems, one may, however, expect that the quality of predictions might be more sensitive to the use of the density-based encoding.




Fig. 2. Example of the density-based encoding when the dimensionality of representation is set to  $N = 4$ .

trainable connections between the hidden and the output layers in the  $\mathbf{W}^{\text{out}}$  matrix is via solving the ridge regression (which is a special case of the Tikhonov regularization) problem, which minimizes the mean square error between predictions (2) and the ground truth. In particular, the activations of the hidden layer  $\mathbf{h}^T$  for each training example are collected together in matrix  $\mathbf{H} \in [M \times N]$ . Matrix  $\mathbf{Y} \in [M \times L]$  stores the corresponding ground-truth classifications using one-hot encodings. Given  $\mathbf{H}$  and  $\mathbf{Y}$ ,  $\mathbf{W}^{\text{out}}$  is calculated as follows:

$$\mathbf{W}^{\text{out}} = (\mathbf{H}^T\mathbf{H} + \lambda\mathbf{I})^{-1}\mathbf{H}^T\mathbf{Y} \quad (3)$$

where  $\mathbf{I}$  denotes an identity matrix of the suitable dimensionality ( $\mathbf{I} \in [N \times N]$ );  $\lambda$  is a hyperparameter (scalar) determining the weight of the regularization part.<sup>4</sup>

### B. Density-Based Encoding of Scalars

The idea of representing scalars as vectors is not new. It has been independently proposed in several areas. The area of stochastic computing [4] is probably the most notable example since the whole idea of the stochastic computing is that it is possible to implement arithmetics on scalars using Boolean operations on vectors (in general, streams) of bits. The rate coding model of neuronal firing used, e.g., in spiking neural networks, is another notable example. Stochastic computing operates with scalars between 0 and 1, which are represented as random bit vectors where the scalar being encoded determines the probability of generating ones. Thus, the density of ones in the obtained bit vector encodes the scalar; hence, such a representation method is called the density-based encoding. Generating random streams is important because the independence of two vectors is a prerequisite for using the Boolean operations to implement the arithmetics on them (e.g., AND for multiplication). Note that, for the proposed approach, no arithmetic operations will be performed with the density-based encodings of scalars. Therefore, the randomness of representations for encoding scalars is not compulsory in this study. In fact, from the simplicity point of view, it is more advantageous to use a structured version of the density-based encoding, which does not require a source of randomness. We will use the structured version of the density-based encoding also known under the name thermometric encoding [9] for the rest of this brief.

The most intuitive way of presenting the concept of the density-based encoding is via visualization. Fig. 2 illustrates all possible values, which could be encoded when the dimensionality of the representation<sup>5</sup> is set to  $N = 4$ . Fig. 2 indicates that using the density-based encoding, it is possible to represent  $N + 1$  different values.

<sup>4</sup>Note that (3) is computationally simpler compared with, e.g., the backprop algorithm, and it is implementable efficiently on CPUs as well as on GPUs.

<sup>5</sup>It will become evident in Section III why the same notation  $N$  as for the number of hidden neurons is used.

The most convenient way of denoting these values is by using integers in the range  $[0, N]$  (nodes on the left in the figure). In this case, in order to obtain the encoding of a given value  $v$ , it is necessary to set  $v$  leftmost positions of the vector to “one” (hashed red nodes in the figure), while the rest of the vector is set to “zero” (filled green nodes). In the case of bipolar representations used in the following, “one” corresponds to  $-1$ , while “zero” corresponds to  $1$ .

Recall, however, that input features are not integers in the range  $[0, N]$ . Instead, it is assumed that a feature  $x_i$  is represented by a real number in the range  $[0, 1]$ . The task is to represent the current value of the feature as a vector  $\mathbf{f} \in [N \times 1]$  using the abovementioned density-based encoding. Since the encoding requires a finite set of values between 0 and  $N$ , real numbers are first discretized using a fixed quantization step, which is determined by  $N$ . Given the current value of the feature, it is quantized to the closest integer as

$$v = \lfloor x_i N \rfloor \quad (4)$$

where  $\lfloor * \rfloor$  denotes rounding to the closest integer. The obtained  $v$  will determine the density-based encoding  $\mathbf{f}$ . The presented procedure allows generating density-based encodings for the whole feature vector  $\mathbf{x}$ . Matrix  $\mathbf{F} \in [N \times K]$ , where  $K$  denotes the number of features, contains the density-based encodings  $\mathbf{f}$  of the current values of  $\mathbf{x}$ .

### C. Hyperdimensional Computing

Hyperdimensional computing [10], [11] also known as vector symbolic architectures is a family of bioinspired methods of representing and manipulating concepts for cognitive architectures and their meanings in a high-dimensional space. Vectors of high (but fixed) dimensionality (denoted as  $N$ ) are the basis for representing information in hyperdimensional computing.<sup>6</sup> The information is distributed across the HD vector’s positions; therefore, HD vectors use distributed representations. Distributed representations [12] are contrary to the localist representations since any subset of the positions can be interpreted. This is very relevant to the density-based encoding introduced in Section II-B since the encoding in  $\mathbf{f}$  is also distributed.

In the scope of this brief paper, columns of  $\mathbf{W}^{\text{in}}$  matrix are interpreted as HD vectors, which are generated randomly. These HD vectors are bipolar ( $\mathbf{W}^{\text{in}} \in \{-1, +1\}^{[N \times K]}$ ) and random with equal probabilities for  $+1$  and  $-1$ . It is worth noting that an important property of high-dimensional spaces is that with an extremely high probability, all random HD vectors are dissimilar to each other (quasi-orthogonal). In order to manipulate HD vectors, hyperdimensional computing defines operations on them. In this brief, we implicitly use only two key operations: binding and bundling.

The binding operation is used to associate two HD vectors together. The result of binding is another HD vector. Here, the result of binding (denoted as  $\mathbf{z}$ ) two vectors  $\mathbf{x}$  and  $\mathbf{y}$  is calculated as follows:  $\mathbf{z} = \mathbf{x} \odot \mathbf{y}$ , where the notation  $\odot$  for the Hadamard product is used to denote the binding operation since this brief uses positionwise multiplication for binding. An important property of the binding operation is that the resultant HD vector  $\mathbf{z}$  is quasi-orthogonal to the HD vectors being bound.

The second operation is called bundling. The bundling operation combines several HD vectors into a single HD vector. Its simplest realization is a positionwise addition. However, when using the positionwise addition, the vector space becomes unlimited; therefore, it is practical to limit the values of the result. This could be achieved

<sup>6</sup>These vectors are referred to as high-dimensional vectors or HD vectors.

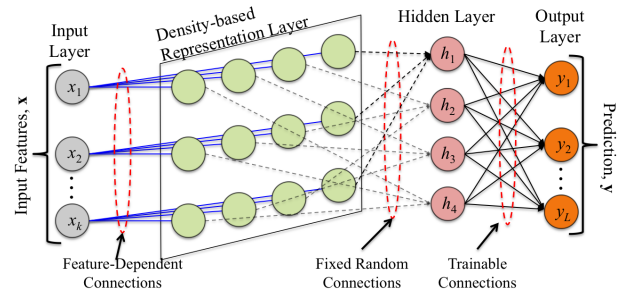


Fig. 3. Architecture of the RVFL that relies on the density-based encoding. In the presented example, the number of hidden neurons and the dimensionality of encoding are set to  $N = 4$ .

with, e.g., a clipping function [denoted as  $f_\kappa(*)$ ]

$$f_\kappa(x) = \begin{cases} -\kappa, & x \leq -\kappa \\ x, & -\kappa < x < \kappa \\ \kappa, & x \geq \kappa. \end{cases} \quad (5)$$

In the clipping function,  $\kappa$  is a configurable threshold parameter. Thus, in this brief, the bundling operation is implemented via positionwise addition limited via the clipping function. For example, the result (denoted as  $\mathbf{a}$ ) of bundling HD vectors  $\mathbf{x}$  and  $\mathbf{y}$  is simply  $\mathbf{a} = f_\kappa(\mathbf{x} + \mathbf{y})$ . In contrast to the binding operation, the resultant HD vector  $\mathbf{a}$  is similar to all bundled HD vectors, which allows, e.g., storing information in HD vectors [13]. For example, we have demonstrated the usefulness of the clipping function for resource-efficient implementations of self-organizing maps [14] and echo state networks [15], [30].

### III. RVFL WITH DENSITY-BASED ENCODINGS

This section presents an architecture of the RVFL utilizing the density-based encoding. The approach is illustrated in Fig. 3. The architecture is intentionally depicted to be as structurally identical to the conventional RVFL (see Fig. 1) as possible. The major difference is that the proposed approach is illustrated with four layers of neurons: input layer ( $\mathbf{x}$ ,  $K$  neurons); density-based representation layer ( $\mathbf{F}$ ,  $N \times K$  neurons); hidden layer ( $\mathbf{h}$ ,  $N$  neurons); and output layer ( $\mathbf{y}$ ,  $L$  neurons). Thus, in contrast to the conventional RVFL, the hidden layer is not connected directly to the input layer. Instead, each input feature is first transformed into a row of neurons storing its density-based encodings. These vectors constitute the density-based representation layer, which, in turn, is connected to the hidden layer. Note also that the input and density-based representation layers are not fully connected. Each neuron in the input layer is only connected to  $N$  neurons in the corresponding row of the next layer. Moreover, these connections (blue lines in Fig. 3) are called “feature-dependent” because the activation of the  $i$ th input neuron  $x_i$  will be quantized to the closest integer  $v$  according to (4); in turn,  $v$  determines the number of the rightmost connections, which transmits  $-1$ , the remaining connections from that neuron transmit  $+1$ . Since each neuron in the density-based representation layer has only one incoming connection, the input activations are projected in the form of the bipolar matrix  $\mathbf{F}$ .

It is also important to mention that the density-based representation and hidden layers are not fully connected. In fact, each neuron in the density-based representation layer has only one outgoing connection. Therefore, the matrix  $\mathbf{W}^{\text{in}}$  describing the fixed random connections to the hidden layer is still  $\mathbf{W}^{\text{in}} \in [N \times K]$ . Moreover, these connections have a clear structure. In Fig. 3, the connections are structured in such a way that each column in  $\mathbf{F}$  is connected to one of the hidden layer neurons. It explains why the number of hidden neurons  $N$



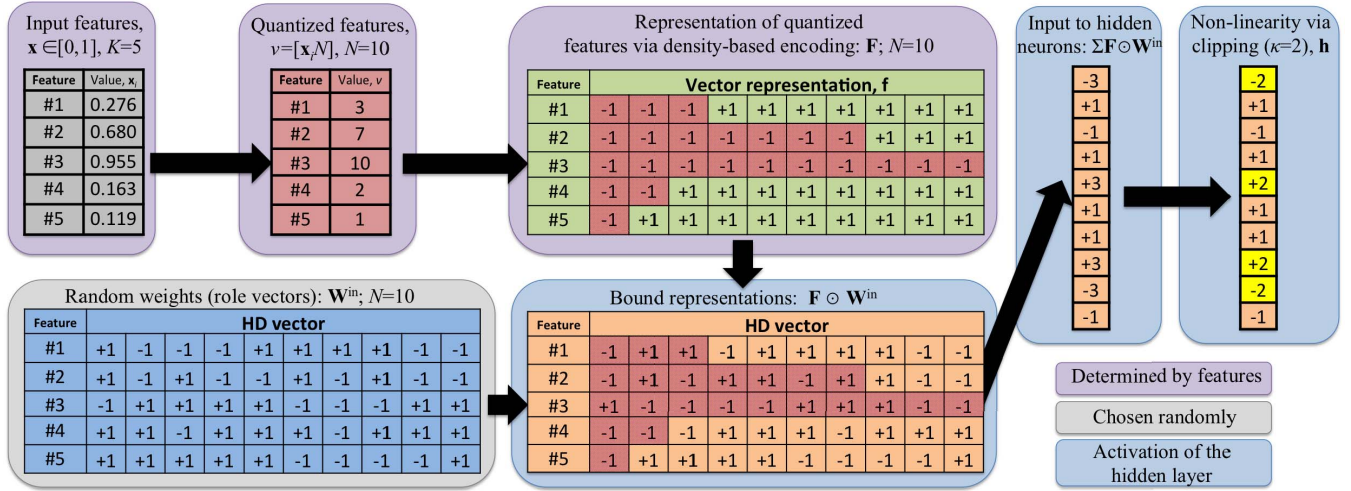


Fig. 4. Example of activating the hidden layer with density-based encodings:  $K = 5$  and  $N = 10$ . Note that  $N$  is set to 10 for visualization purposes only.

also determines the dimensionality of the density-based encoding of features: each hidden neuron has its corresponding column in  $F$  (see Fig. 4). Note that, in Fig. 4,  $N$  is set to 10 only for visualization purposes. In practice, the values of  $N$  are larger.

Similar to the conventional RVFL, the values of  $W^{in}$  are also generated randomly. However, the values are drawn equiprobably from  $\{-1, +1\}$ . Thus, similar to  $F$ ,  $W^{in}$  is also a bipolar matrix. When reflecting to the ideas of hyperdimensional computing,  $W^{in}$  should be interpreted as  $K$   $N$ -dimensional bipolar HD vectors. In other words, each feature is assigned with the corresponding HD vector. Thus, a conceptual intermediate step before getting input values of the hidden neurons is the binding operation between features' HD vectors and their current density-based encoding.

Finally, the proposed approach uses different nonlinear activation function in the hidden layer; the clipping function (5) is used instead of the sigmoid function. The clipping function is characterized by the threshold value  $\kappa$  regulating nonlinear behavior of the neurons and limiting the range of activation values. Summarizing the aforementioned differences, activations of the hidden layer  $h$  are obtained as follows:

$$\mathbf{h} = f_{\kappa} \left( \sum \mathbf{F} \odot \mathbf{W}^{in} \right) \quad (6)$$

where  $\sum$  is a columnwise summation. Note that in contrast to (1), there is no bias term since it has been found empirically that its presence does not improve classification performance. Once the activations of the hidden layer  $h$  are obtained, the rest of the network works in the same way as the conventional RVFL. The predictions in  $\mathbf{y}$  are calculated according to (2).

In order to make operations of the proposed approach more intuitive, Fig. 4 presents a numerical example of acquiring the activations of the hidden layer. First, the input layer with  $K = 5$  neurons sets the values of the current feature vector. These values are quantized to integers in the range  $[0, 10]$  (since  $N = 10$ ). The quantized values determine the neurons of the density-based encoding, which are set to  $-1$  (the rest is  $+1$ ). For example, since the third feature is quantized to  $v = 10$ , all values of its density-based encoding are set to  $-1$ . The bottom left figure shows a randomly generated  $W^{in}$ . Once  $F$  is obtained, we calculate the Hadamard product  $F \odot W^{in}$ , which is denoted as “bound representations” in Fig. 4. The rowwise summation of the resultant matrix represents the input values of the hidden layer. Finally, the clipping function ( $\kappa = 2$  in Fig. 4) is used in the hidden layer to get  $h$ .

Note that due to the way of forming  $F$  and  $W^{in}$ , the input to the hidden layer neurons is always integers in the range  $[-K, K]$ . Moreover, even after the clipping, the activations of neurons are integers in the range  $[-\kappa$  and  $\kappa]$  (practically,  $\kappa < K$ ). Thus, each hidden neuron can be represented using only  $\lceil \log_2(2\kappa + 1) \rceil$  bits of memory. For example, when  $\kappa = 3$ , there are seven unique activations of a neuron, which can be stored with just three bits. Last but not least, it is worth mentioning that, for an efficient implementation, the explicit calculation of  $F$  is redundant. As it could be seen from Fig. 4, the same result as  $F \odot W^{in}$  could be obtained if, for each feature, we use  $v$  as an indicator of which signs should be changed in  $W^{in}$ . As it will be shown in Section IV, these properties give a major advantage over the conventional RVFL for resource-efficient implementation on digital hardware.

Since, in the proposed approach, the part of the network between the hidden and output layers is not modified, the simplest case is to train the readout matrix  $W^{out}$  in the same manner as for the conventional RVFL (see Section II-A). Note that while training the readout matrix  $W^{out}$  according to (3), there is no need to normalize the activation values in  $h$ . Moreover, since the goal is to obtain a very simplistic implementation, it is worth considering alternatives where  $W^{out}$  would contain only integer values in a small limited range. In particular, we have considered three options: quantizing the result of regression (3); using a genetic algorithm (GA) initialized randomly; and using GA initialized with the quantized result of regression. During the search, GA used the cost function for the generalized learning vector quantization [16].

#### IV. PERFORMANCE EVALUATION

In this section, the proposed approach is verified in three scenarios.<sup>7</sup> The first scenario compares it against the conventional RVFL in the case when the weights of the readout matrix are real numbers for both approaches. The second scenario compares the results for the real-valued readout matrix against the considered strategies of obtaining an integer-valued readout matrix. The final scenario compares FPGA implementations of the proposed approach and the finite precision RVFL [8] in the case of a limited energy budget.

<sup>7</sup>The diligent readers are kindly referred to the Supplementary Material that provides an additional experimental evaluation to further justify the proposed approach.




Fig. 5. Cross-validation accuracy of the conventional RVFL against the proposed approach. A point corresponds to a data set.

All reported results<sup>8</sup> are based on 121 real-world classification data sets obtained from the UCI machine learning repository<sup>9</sup> [17]. The considered collection of data sets has been initially analyzed in a large-scale comparison study of different classifiers, and the interested readers are kindly referred to the original work [18] for more details. The only preprocessing step was to normalize features in the range [0, 1]. Finally, the reported accuracies were averaged across five independent initializations.

#### A. Comparison With the Conventional RVFL

First, we compare the conventional RVFL with the proposed approach when computational resources for both approaches are not limited. The search of the hyperparameters has been done according to [18] using the grid search over  $\lambda$  and  $N$  in the case of the conventional RVFL and additionally considering  $\kappa$  for the proposed approach;  $N$  varied in the range [50, 1500] with step 50;  $\lambda$  varied in the range  $2^{[-10,5]}$  with step 1; and  $\kappa$  varied between {1, 3, 7, 15}. The obtained optimal hyperparameters were used to estimate the cross-validation accuracy on all data sets. In order to visualize the results, we rely on the same approach, as reported in [19]. Fig. 5 presents the accuracy of the conventional RVFL against the proposed approach. First, it is important to note that, as expected, the correlation coefficient between the obtained results is high (0.86). Moreover, the average accuracy for the conventional RVFL is 0.76, while that for the proposed approach is 0.80.<sup>10</sup> The difference in accuracy was statistically significant using a 5% significance level according to two-sample hypothesis testing. It is not absolutely intuitive why the proposed approach demonstrates higher accuracy. Nevertheless, one hypothesis is that the quantization for the density-based encoding might provide extra regularization.<sup>11</sup>

#### B. Effect of Quantized Readout Weights

Fig. 6 presents the average accuracy of the proposed approach for three considered strategies of obtaining the readout matrix with

<sup>8</sup>The implementation of the experiments reported in this brief is available online via <https://github.com/denk1e/Density-Encoding-Enables-Resource-Efficient-Randomly-Connected-Neural-Networks>

<sup>9</sup>Available online: <http://persoal.citius.usc.es/manuel.fernandez.delgado/papers/jmlr/data.tar.gz>

<sup>10</sup>In [18], the highest mean accuracy 0.82 was obtained for the random forest.

<sup>11</sup>For the conventional RVFL with quantized inputs, the mean accuracy was 0.753 (0.755 for nonquantized). The correlation coefficient was 0.986. Thus, the improvement cannot be caused barely by the input quantization. Also, for the case when the conventional RVFL additionally used the connections between the input layer and the output layer, there was no significant improvement in the accuracy as the mean accuracy was 0.762.




Fig. 6. Average cross-validation accuracy of the proposed approach for different integer readout strategies.

integer values against the average accuracy from the previous experiment. The considered ranges are symmetric, and the figure indicates only positive boundaries. It is clear that if the result of regression is quantized (dashed line) to very few levels, the accuracy is affected significantly. However, with the increased number of levels, the accuracy approaches the baseline, and it is concluded that 5 bits per weight result in a very close approximation. Refining the quantized result of regression with GA (dash-dotted line) certainly improved the accuracy for a small number of quantization levels, which is in line with the results in [8]. However, using GA for the number of levels larger than six was not beneficial. Random GA initialization (dotted line) decreased the accuracy.

#### C. Performance in the Case of Limited Resources

The third experiment compares FPGA hardware implementations of the proposed approach and the finite precision RVFL in the case of a fixed energy budget per one classification pass. The idea of restricting the energy budget could be seen as an intuitive setup for comparing bounded-optimality [20] of two approaches. Finite precision RVFL [8] with 8 bits per neuron/weight was used since it is more efficient than the conventional RVFL.<sup>12</sup> Following the conclusions from the previous experiment, the resolution of the readout weights of the proposed approach was set to 5 bits. Both approaches were deployed on ZedBoard FPGA, and the energy consumption was estimated with the Xilinx Power Estimator tool. The energy budget was set to 3.2  $\mu$ J to reflect a network with typical parameters.<sup>13</sup> Fig. 7 presents the accuracy of the proposed approach (average 0.73) against the finite precision RVFL (average 0.65). The difference in accuracy was statistically significant using a 5% significance level according to two-sample hypothesis testing. Due to the limited resources, values are lower than in the first experiment; nevertheless, the results are impressive when the performance of our approach is compared with the fixed point RVFL.

## V. RELATED WORK

This section briefly describes the related work. First, the readers generally interested in neural networks, which rely on randomly created connections, are kindly referred to the survey in [6].

<sup>12</sup>When comparing FPGA implementations of the conventional RVFL and the proposed approach (integer readout) for a network where  $K = 16$ ,  $N = 512$ , and  $L = 4$  (median values for the 121 UCI data sets), the proposed approach consumed about 11 times less energy and was 2.5 times faster.

<sup>13</sup>Since, in each data set, the input and output sizes are fixed, the budget was enforced by determining the number of hidden layer neurons within the budget.

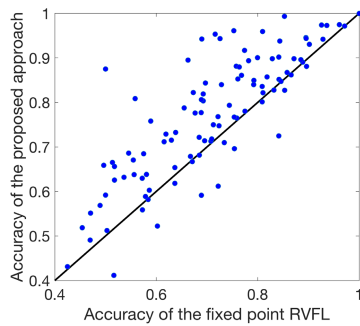


Fig. 7. Cross-validation accuracy of the finite precision RVFL against the proposed approach in case of the fixed energy budget.

### A. Paradigms Used for the Proposed Approach

In order to design a resource-efficient RVFL algorithm, the proposed approach combines the ideas from two areas. These are the density-based encoding from stochastic computing and the binding and bundling operations from hyperdimensional computing. Since both are research fields on their own, here, we only indicate the introductory papers facilitating entrance to the areas. The recent magazine article [4] is probably the most approachable reading for stochastic computing. With respect to hyperdimensional computing, the best starting point is the tutorial-like article by Kanerva [5].

### B. Resource-Efficient RVFL

Recall that even the conventional RVFL networks are considered to be one of the simplest approaches for machine learning. This fact explains why the efforts on pushing the resource-efficiency of RVFL networks to the extreme are rather limited. The most relevant works in this direction are [8] and [21]. Similar to this study, both works use FPGA for hardware experiments. Moreover, both works rely on finite precision implementation for improving the resource-efficiency. The work [21], however, heavily focuses on the process of obtaining the weights of the readout matrix, which is not the case here. The work [8], which focuses on the operational phase, is used here as the baseline for comparison with the proposed approach. However, none of the previous works in the area of RVFL, to the best of our knowledge, have been focusing on using the combination of the density-based encoding with the binding operation. As an important topic for future research, we see the theoretical characterization of the classification performance improvement obtained with the proposed approach. As indicated earlier, one hypothesis is that the quantization and the density-based encoding provide extra regularization. In order to move in this direction, the relevant works are related to a phenomenon of network generalization improvement, e.g., via adding a noise [22] or discretizing quantitative features [23].

### C. Simplification of Neural Networks

Finally, it is worth mentioning that, in recent years, the simplification of computing architectures for neural networks is an important research topic. Notable examples are works [24]–[26], which have been evaluated on convolutional neural networks; work [27] that has introduced networks with ternary activations and work [28] that has introduced networks where all parameters are binary. It is worth mentioning that, in contrast to the bitwise networks [28], the proposed use of the density-based encoding does not require the binarization of the input features, which often worsens the accuracy.

## VI. CONCLUSION

This article proposed a resource-efficient fully integer approach to randomly connected neural networks. The key enabler for efficiently obtaining activations of the hidden neurons is the combination of the representation of input features via the density-based encoding used in the stochastic computing and the use of binding and bundling operations from hyperdimensional computing area. Integer values of the readout matrix could be obtained with a minimal loss in the accuracy, e.g., by simple rounding of the ridge regression solution, which, in turn, could be fine-tuned by the GA. The empirical evaluation was performed on 121 real-world data sets. The proposed approach demonstrated a higher average accuracy than the conventional RVFL networks while being 2.5 times faster and consuming 11 times less energy (typical network on FPGA). Finally, the accuracy of the proposed approach significantly prevailed that of the finite precision RVFL networks when both networks implemented on hardware were constrained to a fixed energy budget.

Despite that this work has focused only on classification tasks, it is worth mentioning that the proposed approach of forming the activations of the hidden layer should be seen as a generic structured representation scheme based on high-dimensional random projections that allow for direct learning of complex nonlinear functions. Therefore, a promising direction for future work is to develop an analytical theory similar to the capacity theory of such representations [13] that would relate the quality of approximations based on the complexity of a nonlinear function, the number of its inputs and outputs, and a number hidden neurons and their resolution.

Last but not least, we conjecture that the density-based encoding will be useful for developing resource-efficient versions of other neural networks. For example, as it has been recently shown in [29] that binarizing initial layers of convolutional neural networks could easily harm their accuracy. We expect that the density-based encoding of input features will solve this issue.

## REFERENCES

- [1] D. Wang and M. Li, "Stochastic configuration networks: Fundamentals and algorithms," *IEEE Trans. Cybern.*, vol. 47, no. 10, pp. 3466–3479, Oct. 2017.
- [2] B. Igel'nik and Y.-H. Pao, "Stochastic choice of basis functions in adaptive function approximation and the functional-link net," *IEEE Trans. Neural Netw.*, vol. 6, no. 6, pp. 1320–1329, Nov. 1995.
- [3] H. D. Nguyen, D. Wang, and G. J. McLachlan, "Randomized mixture models for probability density approximation and estimation," *Inf. Sci.*, vol. 467, pp. 135–148, Oct. 2018.
- [4] A. Alaghi and J. P. Hayes, "Computing with randomness," *IEEE Spectr.*, vol. 55, no. 3, pp. 46–51, Mar. 2018.
- [5] P. Kanerva, "Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors," *Cognit. Comput.*, vol. 1, no. 2, pp. 139–159, Jun. 2009.
- [6] S. Scardapane and D. Wang, "Randomness in neural networks: An overview," *Data Mining Know. Discovery*, vol. 7, pp. 1–18, 2017.
- [7] L. Zhang and P. N. Suganthan, "A comprehensive evaluation of random vector functional link networks," *Inf. Sci.*, vols. 367–368, pp. 1094–1105, Nov. 2016.
- [8] A. Rosato, R. Altilio, and M. Panella, "Finite precision implementation of random vector functional-link networks," in *Proc. 22nd Int. Conf. Digit. Signal Process. (DSP)*, Aug. 2017, pp. 1–5.
- [9] D. A. Rachkovskij, S. V. Slipchenko, E. M. Kussul, and T. N. Baidyk, "Sparse binary distributed encoding of scalars," *J. Autom. Inf. Sci.*, vol. 37, no. 6, pp. 12–23, 2005.
- [10] T. A. Plate, *Holographic Reduced Representations: Distributed Representation for Cognitive Structures*. Stanford, CA, USA: CSLI, 2003.
- [11] S. I. Gallant and T. W. Okaywe, "Representing objects, relations, and sequences," *Neural Comput.*, vol. 25, no. 8, pp. 2038–2078, Aug. 2013.

- [12] G. E. Hinton, J. L. McClelland, and D. E. Rumelhart, "Distributed representations," in *Proc. PDP*, 1986, pp. 77–109.
- [13] E. P. Frady, D. Kleyko, and F. T. Sommer, "A theory of sequence indexing and working memory in recurrent neural networks," *Neural Comput.*, vol. 30, no. 6, pp. 1449–1513, Jun. 2018.
- [14] D. Kleyko, E. Osipov, D. D. Silva, U. Wiklund, and D. Alahakoon, "Integer self-organizing maps for digital hardware," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, Jul. 2019, pp. 1–8.
- [15] D. Kleyko, E. Paxon Frady, and E. Osipov, "Integer echo state networks: Hyperdimensional reservoir computing," 2017, *arXiv:1706.00280*. [Online]. Available: <http://arxiv.org/abs/1706.00280>
- [16] A. Sato and K. Yamada, "Generalized learning vector quantization," in *Proc. NeurIPS*, 1996, pp. 423–429.
- [17] D. Dua and C. Graff. (2019). *UCI Machine Learning Repository*. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [18] M. Fernández-Delgado, E. Cernadas, S. Barro, and D. Amorim, "Do we need hundreds of classifiers to solve real world classification problems?" *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 3133–3181, 2014.
- [19] M. Olson, A. J. Wyner, and R. Berk, "Modern neural networks generalize on small data sets," in *Proc. NeurIPS*, 2018, pp. 1–10.
- [20] S. J. Russell and D. Subramanian, "Provably bounded-optimal agents," *J. Artif. Intell. Res.*, vol. 2, pp. 575–609, May 1995.
- [21] J. M. Martínez-Villena, A. Rosado-Muñoz, and E. Soria-Olivas, "Hardware implementation methods in random vector functional-link networks," *Int. J. Speech Technol.*, vol. 41, no. 1, pp. 184–195, Jul. 2014.
- [22] C. M. Bishop, "Training with noise is equivalent to tikhonov regularization," *Neural Comput.*, vol. 7, no. 1, pp. 108–116, Jan. 1995.
- [23] N. A. Zaidi, Y. Du, and G. I. Webb, "On the effectiveness of discretizing quantitative attributes in linear classifiers," 2017, *arXiv:1701.07114*. [Online]. Available: <http://arxiv.org/abs/1701.07114>
- [24] B. Jacob *et al.*, "Quantization and training of neural networks for efficient Integer-Arithmetic-Only inference," 2017, *arXiv:1712.05877*. [Online]. Available: <http://arxiv.org/abs/1712.05877>
- [25] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, and Y. Bengio, "Binarized neural networks," in *Proc. NeurIPS*, 2016, pp. 1–9.
- [26] I. Hubara, M. Courbariaux, and D. Soudry, "Quantized neural networks: Training neural networks with low precision weights and activations," *J. Mach. Learn. Res.*, vol. 18, pp. 1–30, Jan. 2018.
- [27] H. Alemdar, V. Leroy, A. Prost-Boucle, and F. Petrot, "Ternary neural networks for resource-efficient AI applications," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, May 2017, pp. 2547–2554.
- [28] M. Kim and P. Smaragdakis, "Bitwise neural networks," in *Proc. ICML*, 2015, pp. 1–5.
- [29] A. G. Anderson and C. P. Berg, "The high-dimensional geometry of binary neural networks," in *Proc. ICLR*, 2018, pp. 1–13.
- [30] O. Nepomnyashchiiy, A. Khantimirov, D. Galayko and N. Sirotinina, "Method of Recurrent Neural Network Hardware Implementation," in *Computer Science On-line Conference: Artificial Intelligence and Bioinspired Computational Methods (CSOC)*, Advances in Intelligent Systems and Computing, vol. 1225, pp. 429–437, 2020.