

Handling dropout probability estimation in convolution neural networks using meta-heuristics

Gustavo H. de Rosa¹ · João P. Papa¹  · Xin-S. Yang²

Published online: 23 June 2017
© Springer-Verlag GmbH Germany 2017

Abstract Deep learning-based approaches have been paramount in recent years, mainly due to their outstanding results in several application domains, ranging from face and object recognition to handwritten digit identification. Convolutional neural networks (CNNs) have attracted a considerable attention since they model the intrinsic and complex brain working mechanisms. However, one main shortcoming of such models concerns their overfitting problem, which prevents the network from predicting unseen data effectively. In this paper, we address this problem by means of properly selecting a regularization parameter known as dropout in the context of CNNs using meta-heuristic-driven techniques. As far as we know, this is the first attempt to tackle this issue using this methodology. Additionally, we also take into account a default dropout parameter and a dropout-less CNN for comparison purposes. The results revealed that optimizing dropout-based CNNs is worthwhile, mainly due to the easiness in finding suitable dropout probability values, without needing to set new parameters empirically.

Keywords Convolutional neural networks · Dropout · Meta-heuristic optimization

Communicated by V. Loia.

✉ João P. Papa
papa@fc.unesp.br
Gustavo H. de Rosa
gustavo.rosa@fc.unesp.br
Xin-S. Yang
x.yang@mdx.ac.uk

¹ Department of Computing, São Paulo State University, Bauru, SP 17033-360, Brazil

² School of Science and Technology, Middlesex University, London NW4 4BT, UK

1 Introduction

One of the main computer vision problems concerns how to produce a good representation of the real world, thus allowing machine learning systems to understand these descriptions for further detecting and classifying objects (Bishop 1995). However, the problem still persists when faced with situations where there exist variations of luminosity in the environment, as well as different perspectives in the image acquisition process and problems related to rotation, translation and scale. Therefore, the great question to be answered by researchers in this area concerns how human beings and animals basically learn by looking around (LeCun et al. 2010).

Standard machine learning approaches devise to settle the aforementioned situation by extracting feature vectors, feeding a classifier based on a training set, and thereafter classifying the remaining images. Thus, even though the feature learning problem has received a substantial attention in the last decades, a spotlight has been lit concerning the study of deep learning techniques (LeCun et al. 1998; Arel et al. 2010; Socher et al. 2011; Farabet et al. 2013; Kavukcuoglu et al. 2010). As these methods are based on the hierarchical feature learning, we can establish an analogy with the human visual processing, i.e., with the information being compressed throughout the vision (eyes) and learning (visual cortex) procedure.

Notwithstanding that there exist several deep learning techniques in the literature, one of the most broadly used approaches concerns the convolutional neural networks (CNNs) (LeCun et al. 1998). These neural networks are formed by different stages and architectures, which are responsible for learning different information at each level (e.g., images and signals). The stages comprehend, basically, filtering operations known as convolution using different

masks (kernels), followed by nonlinear operations known as pooling for posterior image subsampling. Each set of these three operations is then executed once again in a new layer, but now with the subsampled image generated as an output of the previous layer, being used as an input to this new layer. Each of these layers creates a new feature vector, which is concatenated at the last layer to compose a high-dimensional feature vector to represent the input image.

Convolutional neural networks have been widely applied in several applications, such as handwritten digit (Simard et al. 2003) and object recognition (LeCun et al. 2004), as well as natural language processing (Collobert and Weston 2008). One can observe that most applications are image processing and computational vision-based, since some CNNs are naturally invariant to translation, rotation, and scaling. However, training deep neural networks with a huge number of parameters can lead to some difficult issues such as overfitting and parameter tuning. The main idea behind the first issue is that during training, units can adapt to the weights drawn from the limited training data, being not that effective when collated with the testing data. Basically, in other words, when we deal with complex information, the units will accommodate themselves during the training step, and “memorize” the data instead of learning, thus resulting in a model that poorly predicts any new or unseen data.

Several attempts have been made to solve the overfitting problem, such as stopping the training as soon as its performance on the validation set starts to drop, or even introducing some types of regularization methods, such as soft-weight sharing (Nowlan and Hinton 1992). On the other hand, a better alternative to address a regularization method would be averaging the predictions of all possible parameter configurations, weighting all the possibilities and checking which one would perform better. Nevertheless, this would require a huge computational effort, thus being only practical for small or simple models (Xiong et al. 2011).

Recently, Srivastava et al. (2014) proposed a technique to overcome these issues, acknowledged as dropout. The pivotal concept about the term “dropout” refers to dropping out units from neural network layers, in other words, by provisionally removing them out of the network, along with all their outgoing and incoming connections. Similarly, Wan et al. (2013) proposed the *DropConnect*, where the authors dropped out connections instead of neurons, which means one can retain a neuron “partially,” since original dropout removes all connections at once. Finally, Iosifidis et al. (2015) presented *DropELM*, in which dropout and DropConnect are combined in the context of Extreme Learning Machines.

Nevertheless, one may find just a few recent works that face the problem of overfitting in CNNs using dropout regularization. For instance, Wu and Gu (2015) introduced a max-pooling dropout method, while Dahl et al. (2013) combined the use of rectified linear unit (ReLU) with dropout.

Although all these techniques have obtained state-of-the-art results in a number of applications, their main drawback concerns fine-tuning the probability of dropping out a neuron, which is usually accomplished by hand.

Considering the drawback of fine-tuning parameters, some recent works attempted to tackle this issue by means of meta-heuristic techniques. For example, Papa et al. (2015a, b, 2016) employed several meta-heuristics to calibrate Bernoulli Restricted Boltzmann Machines, Discriminative Restricted Boltzmann Machines and Deep Belief Networks parameters, respectively. Moreover, Rosa et al. have used meta-heuristics for optimizing CNNs (2015) and Deep Belief Networks parameters (2016).

In this paper, the problem of selecting a proper dropout probability in CNNs is modeled as a meta-heuristic-driven optimization task, in which agents encode the values of the probabilities in a search problem guided by the loss function over the validation set. As far as we are concerned, this is the first work that attempted to address the problem of properly selecting the dropout ratio in such a way. In order to validate the proposed approach, we employed Particle Swarm Optimization (PSO) (Kennedy and Eberhart 2001), since it is a well-known and consistent meta-heuristic optimization technique, as well as Bat Algorithm (BA) (Yang 2010b), Cuckoo Search (CS) (Yang and Deb 2010) and Firefly Algorithm (FA) (Yang 2010a), which are bio-inspired techniques that turned out to be a recent hotbed due to their good effectiveness in several applications.

Therefore, the main contributions of this paper are twofold: (i) to introduce meta-heuristic techniques in the context of dropout regularization in CNNs and (ii) to fill the lack of research regarding dropping out units in CNNs. The remainder of this paper is organized as follows. Sections 2 and 3 present some theoretical background with respect to CNNs and dropout techniques, respectively. Section 4 discusses the methodology employed in this work, and Sect. 5 presents the experiments. Finally, Sect. 6 states conclusions and future works.

2 Convolutional neural networks

Convolutional neural networks can be seen as a representation of a bigger class of models based on the Hubel’s and Wiesel’s architecture, which was presented in a seminal study in 1962 concerning the primary cortex of cats. This research has identified, basically, two kinds of cells: (i) *simple cells*, which possess an analogous duty to the filter bank step, and (ii) the *complex cells*, which perform a similar job to the CNN sampling step.

The first model that simulated a computer-based convolutional neural network was the well-known “Neocognitron” (Fukushima and Miyake 1982), which implemented an unsu-

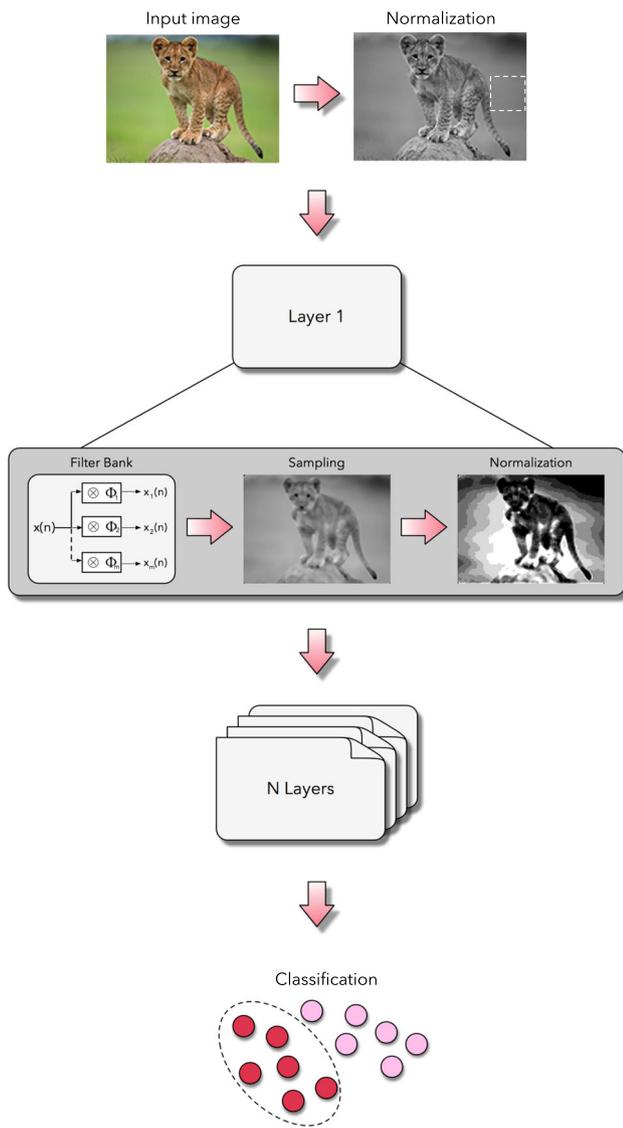


Fig. 1 A typical convolutional neural network architecture

pervised training algorithm during the filter bank step, followed by a supervised training algorithm applied in the last layer. Later on, [LeCun et al. \(1989\)](#) simplified this architecture by proposing the use of the backpropagation algorithm to train the network in a supervised way. Thus, several applications that used CNN emerged in the subsequent decades.

Basically, a CNN can be understood as an N -layered data processing sequence. Thereby, given an input image,¹ a CNN essentially extracts a high-level representation of it, called *multispectral image*, whose pixel attributes are concatenated in a feature vector for later application of pattern recognition techniques. Figure 1 introduces the naïve architecture of a convolutional neural network.

¹ The same procedure can be extended to signal processing-based applications.

As aforementioned, each CNN layer is often composed of three operations, being the first one a convolution with a filter bank, followed by a sampling phase and then by a normalization step. As one can observe in Fig. 1, there is still a possibility of a normalization operation in the beginning of the whole process. The next sections describe in more details each of these steps.

2.1 Filter bank

Let $\hat{I} = (D_I, \mathbf{I})$ be a multispectral image such that $D_I \in n \times n$ is the image domain, and $\mathbf{I} = \{I_1(p), I_2(p), \dots, I_m(p)\}$ corresponds to a pixel $p = (x_p, y_p) \in D_I$, and m stands for the number of bands. When \hat{I} is a grayscale image, for instance, we have that $m = 1$ and $\hat{I} = (D_I, I)$.

Let $\phi = (\mathcal{A}, W)$ be a filter with weights $W(q)$ associated with every *pixel* $q \in \mathcal{A}(p)$, where $\mathcal{A}(p)$ denotes a mask of size $L_A \times L_A$, centered at p , and $q \in \mathcal{A}(p)$ if, and only if, $\max\{|x_q - x_p|, |y_q - y_p|\} \leq (L_A - 1)/2$. In case of multispectral filters, their weights can be depicted as vectors $\mathbf{W}_i(q) = \{w_{i,1}(q), w_{i,2}(q), \dots, w_{i,m}(q)\}$ for each filter i , and a multispectral filter bank can be then defined as $\phi = \{\phi_1, \phi_2, \dots, \phi_n\}$, where $\phi_i = (\mathcal{A}, \mathbf{W}_i)$, $i = \{1, 2, \dots, n\}$.

Thus, the convolution between an input image \hat{I} and a filter ϕ_i generates the band i of the filtered image $\hat{J} = (D_J, \mathbf{J})$, where $D_J \in D_I$ and $\mathbf{J} = \{J_1(p), J_2(p), \dots, J_n(p)\}$, $\forall p \in D_J$:

$$J_i(p) = \sum_{\forall q \in \mathcal{A}(p)} \mathbf{I}(q) \otimes \mathbf{W}_i(q), \tag{1}$$

where \otimes denotes the convolution operator. The weights of ϕ_i are usually generated from an uniform distribution, i.e., $U(0, 1)$, and afterward normalized with mean zero and unitary norm.

2.2 Sampling

This operation is extremely important for a CNN, which provides translational invariance to the extracted features. Let $\mathcal{B}(p)$ be the sampling area of size $L_B \times L_B$ centered at p . Additionally, let $D_K = D_J/s$ be a regular sampling operation every s pixels. Therefore, the resulting sampling operation in the image $\hat{K} = (D_K, \mathbf{K})$ is defined as follows:

$$K_i(p) = \sqrt{\alpha \sum_{\forall q \in \mathcal{B}(p)} J_i(q)^\alpha}, \tag{2}$$

where $p \in D_K$ denotes every pixel of the new image, $i = \{1, 2, \dots, n^2\}$, and α stands for the stride parameter, controlling the downsampling factor of the operation.

2.3 Normalization

The last operation of a CNN is its normalization, which is a widely employed mechanism in order to enhance its performance (Cox and Pinto 2011). This operation is based on the apparatus found on cortical neurons (Geisler and Albrecht 1992), being also defined under a squared-area $\mathcal{C}(p)$ of size $L_C \times L_C$ centered at pixel p , such as:

$$O_i(p) = \frac{K_i(p)}{\sum_{j=1}^n \sum_{\forall q \in \mathcal{C}(p)} K_j(q) K_i(q)}. \quad (3)$$

Thus, the above operation is accomplished for each pixel $p \in D_O \subset D_k$ of the resulting image $\hat{O} = (D_O, \mathbf{O})$.

3 Dropout regularization

Considering the aforementioned CNN model, a dropout-based CNN can be formulated as a new CNN layer. In this new formulation, \mathbf{r} stands for the activation or dropout of the M neurons in a specific layer, where each variable r_j contains the value 1 (one) with probability $1 - p$, independent of other variables $r_i, i \neq j$. If r_j equals to 1 (one), the unit h_j is withheld, otherwise it is dropped from the network together with its connections.

Notice that probability p is independent of other units and \mathbf{r} is sampled directly from a Bernoulli distribution (Srivastava et al. 2014), being resampled for every mini-batch during the learning process. Equation (4) describes this distribution:

$$r_j \sim \text{Bernoulli}(p), \forall j = \{1, 2, \dots, M\}. \quad (4)$$

Therefore, during the training phase, let $y^{(L)}$ denote the vector of outputs at a layer L . The new vector of outputs $\tilde{y}^{(L)}$ considering dropout is formulated by Eq. (5):

$$\tilde{y}^{(L)} = r y^{(L)}. \quad (5)$$

Finally, in the testing step, the weight matrix W needs to be scaled with ratio p in order to average all the 2^M possible dropped-out networks. This is the greatest contribution of the dropout regularization, as it only needs to test a single network. Equation (6) is responsible to illustrate this process.

$$W_{\text{test}}^{(L)} = p \mathbf{W}^{(L)}, \quad (6)$$

where $\mathbf{W}^{(L)}$ stands for the weight matrix at layer L .

4 Methodology

In this section, we present the methodology used to properly select the dropout ratio using convolutional neural networks, as well as the employed datasets.

4.1 Modeling dropout-based CNN parameter optimization

We propose to model the problem of selecting suitable dropout parameters considering CNN in the task of image classification. The learning step of a CNN has four parameters: the learning rate η , penalty parameter (momentum) α , weight decay λ , and the dropout ratio p . As we are interested in fine-tuning the dropout parameter only, we fixed the 3-tuple (η, α, λ) and we played with p in order to minimize the loss function of the classified images over a validation set. After that, the selected parameter is thus applied to classify the unseen images of the test set.

4.2 Datasets

In regard to the parameter optimization experiment, we employed four datasets, as described below:

- MNIST dataset:² it is composed of images of handwritten digits. The original version contains a training set with 60,000 images from digits “0”–“9,” as well as a test set with 10,000 images.
- Semeion handwritten digits dataset:³ set of 1593 images from handwritten digits “0”–“9” written in two ways: the first time in a normal way (accurately) and the second time in a fast way (no accuracy). In the end, they were stretched with resolution of 16×16 in a grayscale of 256 values, and then, each pixel was binarized.
- USPS dataset:⁴ numeric data obtained from the scanning of handwritten digits from envelopes by the US Postal Service. It is composed of 7291 training images and a test set with 2007 images. All the original digits were deslanted and size-normalized, resulting in 16×16 grayscale images.
- CIFAR-10 dataset:⁵ it is a subset image database from the “80 million tiny images” dataset, collected by Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Composed of 60,000 32×32 color images in 10 classes, with 6000 images per class. It is also divided into five training batches and one test batch, each one containing 10,000

² <http://yann.lecun.com/exdb/mnist/>.

³ <https://archive.ics.uci.edu/ml/datasets/Semeion+Handwritten+Digit>.

⁴ <http://statweb.stanford.edu/tibs/ElemStatLearn/datasets/zip.info.txt>.

⁵ <http://www.cs.toronto.edu/kriz/cifar.html>.

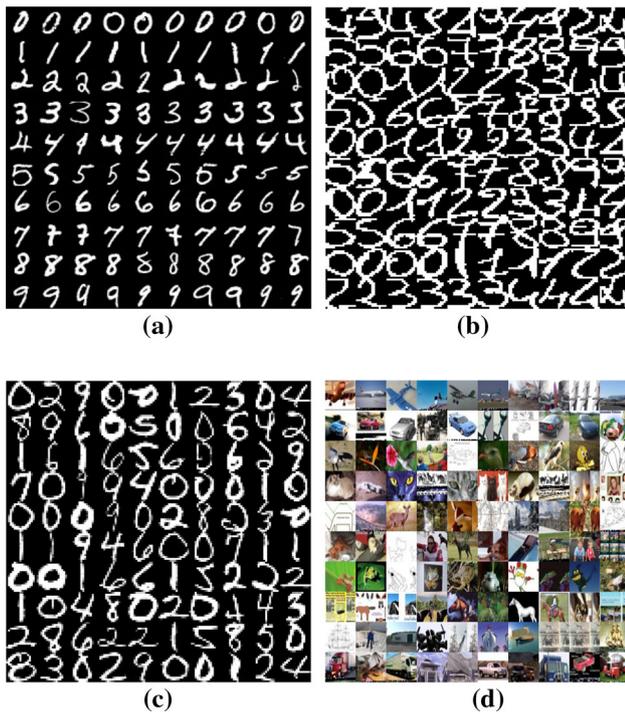


Fig. 2 Some training examples from: **a** MNIST, **b** Semeion, **c** USPS and **d** CIFAR-10 datasets

images. Therefore, we have 50,000 images for training purposes and 10,000 for testing duties.

Figure 2 displays some training examples from the above datasets.

4.3 CNN architectures

In regard to the source code, we used our own optimization library LibOPT⁶ and the well-known Caffe library⁷ (Jia et al. 2014), which is developed under General-Purpose computing on Graphics Processor Units (GPGPU) platform, thus providing more efficient implementations concerning CNNs. Also, in order to integrate LibOPT with Caffe, we developed a new library called LibOPT4Caffe.⁸ Additionally, we considered two different CNN architectures to provide a deeper experimental analysis: one for CIFAR-10 dataset and another for MNIST, Semeion and USPS datasets. These architectures are the original ones proposed by Caffe examples, just with an extra dropout layer. Also, regarding Semeion and USPS datasets, we have used a kernel size of 3 instead of 5 for convolution layers due to the lower resolution of these datasets. Figure 3 illustrates the architectures used in this work. Note

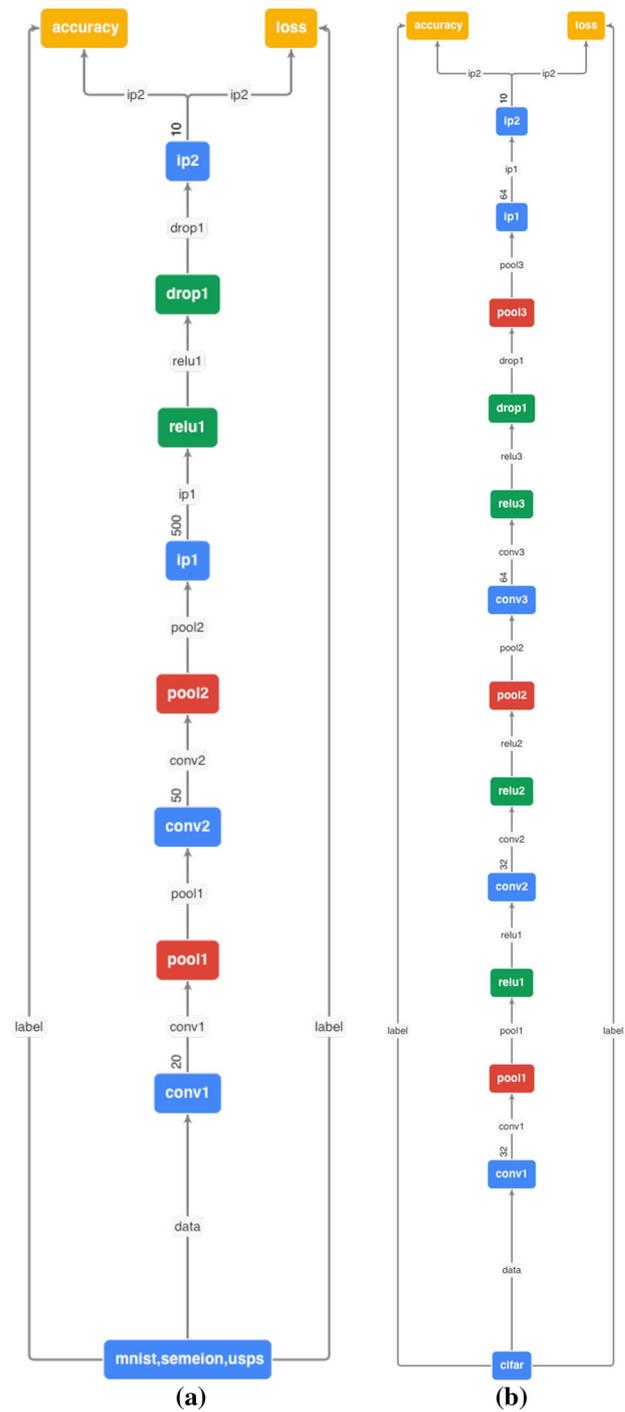


Fig. 3 CNN architectures for: **a** MNIST, Semeion and USPS, and **b** CIFAR-10 datasets. The picture was generated with the following tool: http://yanglei.me/gen_proto/

that “conv” stands for the convolution layer, “pool” for pooling, “relu” for rectified linear unit, “drop” for dropout and “ip” for inner product, which stands for a multiplication in the vector space.

⁶ <https://github.com/jppbsi/LibOPT>.

⁷ <http://caffe.berkeleyvision.org>.

⁸ <https://github.com/gugarosa/LibOPT4Caffe>.

4.4 Experimental setup

In this work, we compared four distinct meta-heuristic optimization techniques against the default dropout ratio given by Caffe, and the standard Caffe architecture without dropout. The main idea to compare meta-heuristics against the default dropout ratio is to check whether it is better to perform or not an optimization prior to the experiments. In order to provide a statistical analysis by means of Wilcoxon signed-rank test (Wilcoxon 1945), we conducted a cross-validation with 20 runs. As we working with a one-dimensional search space, there is no need to employ a vastly number of agents nor iterations. Thus, we employed 7 agents over 10 iterations for convergence considering all techniques. Table 1 presents the native parameter configuration for each optimization technique. Notice these values have been empirically set.

All datasets were split into training, validation, and testing sets. Table 2 describes the amount of images, as well as the number of the batch size (parenthesis) employed for each set.

Finally, we have set each CNN parameter according to Table 3. Note that all these parameters are the default ones provided by Caffe library. The only exception holds for Semeion dataset due to its small amount of images, being trained in a slower pace with a learning rate ten times smaller, i.e., 0.001.

Table 1 Meta-heuristic parameter configuration

Technique	Parameters
BA	$f_{\min} = 0, f_{\max} = 2, A = 0.5, \text{rand} = 0.5$
CS	$\beta = 1.5, p = 0.25, \alpha = 0.8$
FA	$\gamma = 1.0, \beta_0 = 1.0, \alpha = 0.2$
PSO	$c_1 = 1.7, c_2 = 1.7, w = 0.7$

Table 2 Dataset configuration

Dataset	# Training set	# Validation set	# Testing set
CIFAR-10	20,000 (100)	30,000 (100)	10,000 (100)
MNIST	20,000 (64)	40,000 (100)	10,000 (100)
Semeion	200 (2)	400 (400)	993 (993)
USPS	2406 (32)	4885 (977)	2007 (2007)

Table 3 CNN parameter configuration

Dataset	η	α	λ	p	# Iterations
CIFAR-10	0.001	0.9	0.004	[0, 1]	4000
MNIST	0.01	0.9	0.0005	[0, 1]	10,000
Semeion	0.001	0.9	0.0005	[0, 1]	10,000
USPS	0.01	0.9	0.0005	[0, 1]	10,000

5 Experimental results

This section aims at presenting the experimental results concerning CNN dropout parameter fine-tuning. We compared four optimization methods against the default dropout-less and the default dropout ratio provided by Caffe. The most accurate results, according to Wilcoxon signed-rank test, are in bold.

Tables 4, 5, 6, and 7 present the average accuracies and average hyper-parameters found over MNIST, Semeion, USPS and CIFAR-10 datasets, respectively. As one can observe, the accuracy rates between different datasets are inconsistent, mainly due to their distinct natures, i.e., amount of images and features, odd acquisition processes, among others.

Table 4 Average accuracies and average hyper-parameters over the test set considering MNIST dataset

	Accuracy (%)	Hyper-parameters			
		η	α	λ	p
Caffe	99.07	0.01	0.9	0.0005	0
Dropout Caffe	99.18	0.01	0.9	0.0005	0.5
BA	99.13	0.01	0.9	0.0005	0.4988
CS	99.14	0.01	0.9	0.0005	0.4883
FA	99.16	0.01	0.9	0.0005	0.4630
PSO	99.17	0.01	0.9	0.0005	0.4559

Table 5 Average accuracies and average hyper-parameters over the test set considering Semeion dataset

	Accuracy (%)	Hyper-parameters			
		η	α	λ	p
Caffe	97.62	0.001	0.9	0.0005	0
Dropout Caffe	98.14	0.001	0.9	0.0005	0.5
BA	98.30	0.001	0.9	0.0005	0.7532
CS	98.19	0.001	0.9	0.0005	0.7546
FA	98.23	0.001	0.9	0.0005	0.7786
PSO	97.66	0.001	0.9	0.0005	0.8263

Table 6 Average accuracies and average hyper-parameters over the test set considering USPS dataset

	Accuracy (%)	Hyper-parameters			
		η	α	λ	p
Caffe	95.80	0.01	0.9	0.0005	0
Dropout Caffe	96.21	0.01	0.9	0.0005	0.5
BA	96.43	0.01	0.9	0.0005	0.7688
CS	96.28	0.01	0.9	0.0005	0.6870
FA	96.38	0.01	0.9	0.0005	0.7319
PSO	96.37	0.01	0.9	0.0005	0.8031

Table 7 Average accuracies and average hyper-parameters over the test set considering CIFAR-10 dataset

	Accuracy (%)	Hyper-parameters			
		η	α	λ	p
Caffe	71.47	0.001	0.9	0.004	0
Dropout Caffe	72.08	0.001	0.9	0.004	0.5
BA	71.43	0.001	0.9	0.004	0.6430
CS	71.16	0.001	0.9	0.004	0.6270
FA	71.52	0.001	0.9	0.004	0.6629
PSO	71.55	0.001	0.9	0.004	0.6655

Regarding MNIST dataset, it is possible to state that a dropout network was once again better than the standard one. All the meta-heuristic techniques, except BA, were able to find the best results among the standard dropout ratio provided by Caffe, even using a smaller probability for dropping off the units. If we take a closer look at Fig. 4, it is possible to observe that after 4000 epochs, the no-dropout network stalls at around 0.01% of error, while the dropout networks still introduce some kind of noise, preventing the network

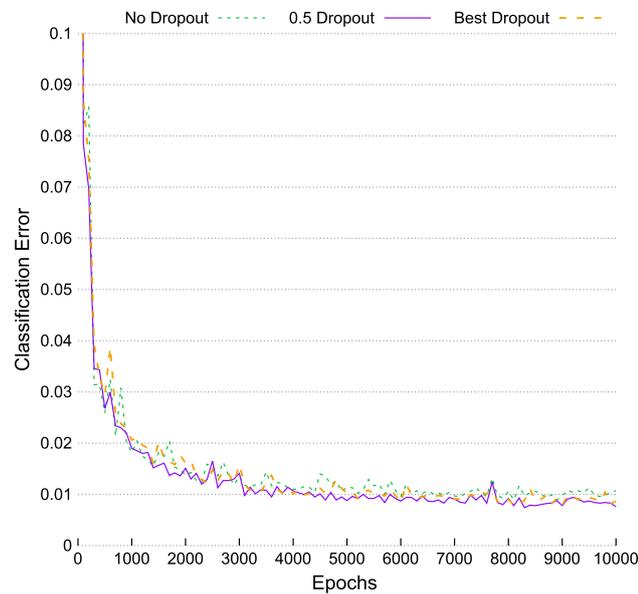


Fig. 4 Classification error over MNIST testing set

from overfitting and improving its classification rate. Note that “Best dropout” concerns the p value found in Table 8.

Considering Semeion dataset, once again dropout was capable of obtaining the best results. Now, all meta-heuristic techniques were also able to obtain the best results, together with our baseline, i.e., dropout by Caffe being BA the one which achieved the highest accuracy over the test set. However, the statistical evaluation pointed out all dropout-based approaches can be considered similar to each other, although the lowest results were obtained by the baseline provided by Caffe. An interesting point concerns the values found for the probability parameter p by the meta-heuristic techniques, which were similar to each other, though being different from the one obtained by the baseline.

Glancing Fig. 5, it is possible to observe the best dropout network was the worst one at the beginning of the epochs. After around 5000 epochs, it starts improving its performance and overcomes the other networks. This is mainly due to the limited amount of images from Semeion dataset, which leads to a significant overfitting and stalls the training after a certain number of epochs.

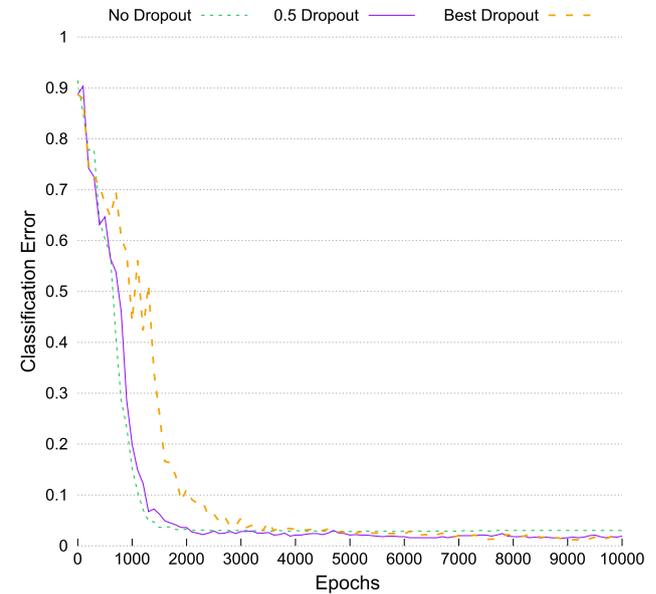


Fig. 5 Classification error over Semeion testing set

Table 8 Average accuracies and best hyper-parameters over the test set considering all datasets

	Accuracy (%)	Best Hyper-parameters				Technique
		η	α	λ	p	
CIFAR-10	71.69	0.001	0.9	0.004	0.6177	BA
MNIST	99.12	0.01	0.9	0.0005	0.5100	PSO
Semeion	98.51	0.001	0.9	0.0005	0.8037	BA
USPS	96.43	0.01	0.9	0.0005	0.8284	FA

BA, FA, and PSO were able to obtain the best results considering USPS dataset. One can observe that their accuracy rates were slightly better than the standard dropout-less and dropout ratio provided by Caffe. As the USPS dataset poses a greater challenge than Semeion one, but still lacks a huge amount of images for deep learning, we can observe in Fig. 6 the search for the optimum dropout ratio performed nicely since the initial epochs. The best dropout network was able to provide a lower classification error almost all the times.

As one can observe, only the default dropout ratio ($p = 0.5$) played a big role on CIFAR-10 dataset, while all the meta-heuristic techniques failed in finding the best parameter for this dataset. However, it is still important to highlight that p was found over the validation set, which is less likely to burden the classification rates. Nevertheless, as already expected, dropout was capable of slightly improving the recognition rate of this architecture. Figure 7 illustrates the classification error over the testing set during all training epochs.

Comprehending our initial experiments, each dataset had a best accuracy rate value for a particular dropout probability p . As one can observe in Table 8, BA, PSO, BA, and FA were the meta-heuristic techniques which found the best p over CIFAR-10, MNIST, Semeion and USPS datasets, respectively. Therefore, we performed another round of experiments using these values. Additionally, Table 8 presents, for each dataset, the average accuracies using the best hyperparameter p . It is important to highlight that, for Semeion dataset, BA was capable of finding a more suitable dropout probability, leading to a significantly higher classification rate than the one observed in Table 5.

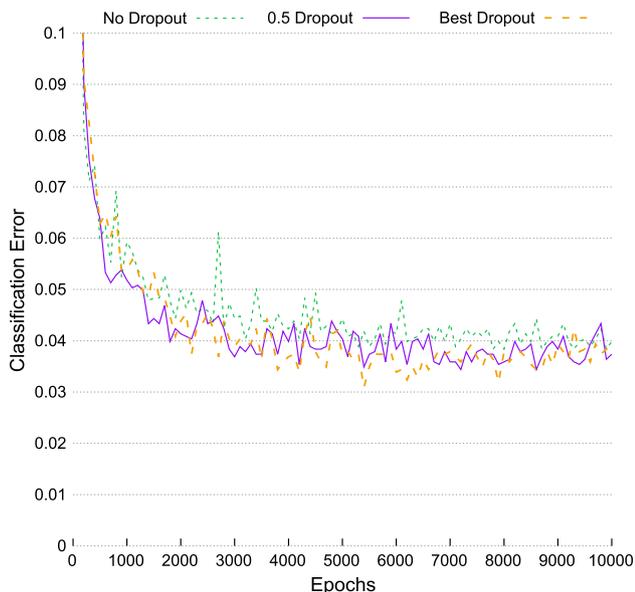


Fig. 6 Classification error over USPS testing set

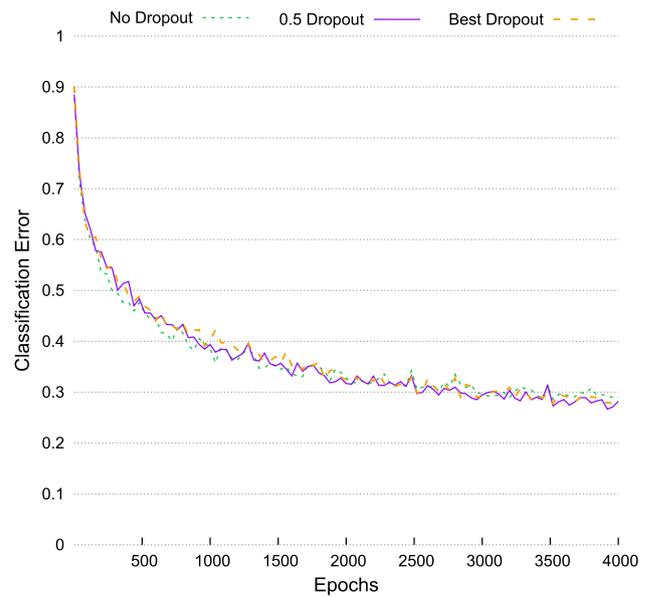


Fig. 7 Classification error over CIFAR-10 testing set

Table 9 Number of CNN evaluations for each technique

Technique	# Calls
Caffe	1
Dropout Caffe	1
BA	78
CS	78
FA	71
PSO	78

Nevertheless, the main drawback concerning meta-heuristic techniques consists in its longer time to attain the final output. As we are evaluating a distinct CNN every time we need to update our particle's fitness, this procedure will take a longer time to converge and find a suitable dropout parameter. BA, CS, and PSO are composed by m initial evaluations + (m agents \times t iterations) + final evaluation with best parameter found, while FA does not have the m initial evaluations. As we are working with $m = 7$ and $t = 10$, Table 9 describes the number of calls to the CNN learning procedure to give us an idea about the computational burden of each technique.

6 Conclusions

In this paper, we dealt with the problem of proper selecting dropout parameters concerning CNNs by means of Particle Swarm Optimization, Bat Algorithm, Cuckoo Search, and Firefly Algorithm. The experiments were carried out over four public datasets in the context of image classification.

The experimental section comprised different CNN architectures, as well as images with different resolutions and

distinct training set sizes. The results obtained by the meta-heuristic-based dropout CNNs were compared against a standard dropout ratio and a dropout-less network, and showed to be very promising, since meta-heuristic CNNs were able to obtain the suitable dropout parameters in almost all datasets. On the other hand, such task requires a higher computational load than no-optimized CNNs, as each particle's fitness update needs to be evaluated under a CNN architecture, thus taking a longer time to find a suitable dropout parameter and achieve its output. Since the meta-heuristic-based techniques use a fitness function based on the CNN's outputs, each iteration of the optimization process requires a full training of the CNN, which turns out to be a costly process. It is usually expected that more iterations and agents would provide better results, but such process comes at the price of a higher computational burden.

In regard to future works, we intend to investigate the proposed approach with other meta-heuristic techniques, as well as to select feasible parameters using meta-heuristics for Dropconnect regularization in the context of CNNs.

Acknowledgements This study was funded by FAPESP (Grants No. #2015/25739-4, #2014/12236-1 and #2014/16250-9) and by CNPq (Grant No. #306166/2014-3).

Compliance with ethical standards

Conflict of interest All authors declare that they have no conflict of interest.

Human participants or animals Notice this article does not contain any studies with human participants or animals performed by any of the authors.

References

- Arel I, Rose D, Karnowski T (2010) Deep machine learning—a new frontier in artificial intelligence research [research frontier]. *IEEE Comput Intell Mag* 5(4):13–18
- Bishop C (1995) *Neural networks for pattern recognition*. Oxford University Press, Oxford
- Collobert R, Weston J (2008) A unified architecture for natural language processing: deep neural networks with multitask learning. In: *Proceedings of the 25th international conference on machine learning*. ACM, pp 160–167
- Cox D, Pinto N (2011) Beyond simple features: a large-scale feature search approach to unconstrained face recognition. In: *Proceedings of the IEEE international conference on automatic face gesture recognition and workshops*, pp 8–15
- Dahl GE, Sainath TN, Hinton GE (2013) Improving deep neural networks for IVCSR using rectified linear units and dropout. In: *2013 IEEE international conference on acoustics, speech and signal processing*. IEEE, pp 8609–8613
- Farabet C, Couprie C, Najman L, LeCun Y (2013) Learning hierarchical features for scene labeling. *IEEE Trans Pattern Anal Mach Intell* 35:1915–1929
- Fukushima K, Miyake S (1982) Neocognitron: a new algorithm for pattern recognition tolerant of deformations and shifts in position. *Pattern Recognit* 15(6):455–469
- Geisler W, Albrecht D (1992) Cortical neurons: isolation of contrast gain control. *Vis Res* 32(8):1409–1410
- Iosifidis A, Tefas A, Pitas I (2015) Dropelm. *Neurocomputing* 162(C):57–66
- Jia Y, Shelhamer E, Donahue J, Karayev S, Long J, Girshick R, Guadarrama S, Darrell T (2014) Caffe: convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*
- Kavukcuoglu K, Sermanet P, Jan Boureau Y, Gregor K, Mathieu M, Cun YL (2010) Learning convolutional feature hierarchies for visual recognition. In: *Advances in neural information processing systems*, vol 23. Curran Associates, Inc, pp 1090–1098
- Kennedy J, Eberhart R (2001) *Swarm intelligence*. M. Kaufman, Los Altos
- LeCun Y, Boser B, Denker JS, Henderson D, Howard RE, Hubbard W, Jackel LD (1989) Backpropagation applied to handwritten zip code recognition. *Neural Comput* 1(4):541–551
- LeCun Y, Bottou L, Bengio Y, Haffner P (1998) Gradient-based learning applied to document recognition. *Proc IEEE* 86(11):2278–2324
- LeCun Y, Huang FJ, Bottou L (2004) Learning methods for generic object recognition with invariance to pose and lighting. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, vol 2, pp 97–104
- LeCun Y, Kavukcuoglu K, Farabet C (2010) Convolutional networks and applications in vision. In: *Proceedings of 2010 IEEE international symposium on circuits and systems*, pp 253–256
- Nowlan SJ, Hinton GE (1992) Simplifying neural networks by soft weight-sharing. *Neural Comput* 4(4):473–493. doi:10.1162/neco.1992.4.4.473
- Papa JP, Rosa GH, Costa KAP, Marana AN, Scheirer W, Cox DD (2015a) On the model selection of Bernoulli restricted Boltzmann machines through harmony search. In: *Proceedings of the genetic and evolutionary computation conference*. ACM, pp 1449–1450
- Papa JP, Rosa GH, Marana AN, Scheirer W, Cox DD (2015b) Model selection for discriminative restricted Boltzmann machines through meta-heuristic techniques. *J Comput Sci* 9:14–18
- Papa JP, Scheirer W, Cox DD (2016) Fine-tuning deep belief networks using harmony search. *Appl Soft Comput* 46:875–885
- Rosa G, Papa J, Marana A, Scheirer W, Cox D (2015) Fine-tuning convolutional neural networks using harmony search. In: *Progress in pattern recognition, image analysis, computer vision, and applications: 20th Iberoamerican congress, proceedings*. Springer International Publishing, pp 683–690
- Rosa G, Papa J, Costa K, Passos L, Pereira C, Yang XS (2016) Learning parameters in deep belief networks through firefly algorithm. In: *Artificial neural networks in pattern recognition: 7th IAPR TC3 workshop, proceedings*. Springer International Publishing, pp 138–149
- Simard PY, Steinkraus D, Platt JC (2003) Best practices for convolutional neural networks applied to visual document analysis. In: *Proceedings of the seventh international conference on document analysis and recognition*. IEEE Computer Society, pp 958–963
- Socher R, Huang EH, Pennin J, Manning CD, Ng AY (2011) Dynamic pooling and unfolding recursive autoencoders for paraphrase detection. In: *Advances in neural information processing systems*, vol 24. Curran Associates, Inc, pp 801–809
- Srivastava N, Hinton GE, Krizhevsky A, Sutskever I, Salakhutdinov R (2014) Dropout: a simple way to prevent neural networks from overfitting. *J Mach Learn Res* 15(1):1929–1958
- Wan L, Zeiler M, Zhang S, LeCun Y, Fergus R (2013) Regularization of neural networks using dropconnect. In: *International conference on machine learning, ICML '13*
- Wilcoxon F (1945) Individual comparisons by ranking methods. *Biom Bull* 1(6):80–83

- Wu H, Gu X (2015) Towards dropout training for convolutional neural networks. *Neural Netw* 71:1–10
- Xiong HY, Barash Y, Frey BJ (2011) Bayesian prediction of tissue-regulated splicing using rna sequence and cellular context. *Bioinformatics* 27(18):2554–2562. doi:[10.1093/bioinformatics/btr444](https://doi.org/10.1093/bioinformatics/btr444)
- Yang XS (2010a) Firefly algorithm, stochastic test functions and design optimisation. *Int J Bio-Inspired Comput* 2(2):78–84
- Yang XS (2010b) A new metaheuristic bat-inspired algorithm. In: *Nature inspired cooperative strategies for optimization (NICSO 2010)*. Springer, pp 65–74
- Yang XS, Deb S (2010) Engineering optimisation by cuckoo search. *Int J Math Model Numer Optim* 1(4):330–343