

Convolutional Channel-Wise Competitive Learning for the Forward-Forward Algorithm

Andreas Papachristodoulou, Christos Kyrkou, Stelios Timotheou, Theodoris Theodoridis

KIOS Research and Innovation Center of Excellence,
Department of Electrical and Computer Engineering, University of Cyprus
{papachristodoulou.andreas, kyrkou.christos, stimo, ttheodoridis}@ucy.ac.cy

Abstract

The Forward-Forward (FF) Algorithm has been recently proposed to alleviate the issues of backpropagation (BP) commonly used to train deep neural networks. However, its current formulation exhibits limitations such as the generation of negative data, slower convergence, and inadequate performance on complex tasks. In this paper we take the main ideas of FF and improve them by leveraging channel-wise competitive learning in the context of convolutional neural networks for image classification tasks. A layer-wise loss function is introduced that promotes competitive learning and eliminates the need for negative data construction. To enhance both the learning of compositional features and feature space partitioning, a channel-wise feature separator and extractor block is proposed that complements the competitive learning process. Our method outperforms recent FF-based models on image classification tasks, achieving testing errors of 0.58%, 7.69%, 21.89%, and 48.77% on MNIST, Fashion-MNIST, CIFAR-10 and CIFAR-100 respectively. Our approach bridges the performance gap between FF learning and BP methods, indicating the potential of our proposed approach to learn useful representations in a layer-wise modular fashion, enabling more efficient and flexible learning. Our source code and supplementary material are available at <https://github.com/andreaspapac/CwComp>.

Introduction

Artificial Neural Networks (ANNs) have emerged as a ubiquitous tool for tackling a variety of machine intelligence tasks. One of the most popular training algorithms for ANNs, Backpropagation (BP), employs the chain rule of differentiation for backward propagation of error gradients, applicable to any differentiable loss function. With its flexibility to accommodate complex architectures involving multiple layers and non-linear activation functions, BP can manage large datasets and high-dimensional input spaces.

However, BP is not without its limitations. Its assumption of perfect knowledge of forward pass computations contributes to the 'black-box' nature of models, challenging interpretability and transparency. Additionally, its task-specific feature learning limits model generalizability, requiring extensive labeled data and increasing susceptibility to adversarial attacks (Moraitis et al. 2022; Madry et al.

2018). Deep neural networks, inherently susceptible to issues like local minima and gradient problems due to non-convexity, face exacerbated challenges with BP's opacity, prompting the need for additional mitigation strategies (Lawrence and Giles 2000; Caruana, Lawrence, and Giles 2000; Tan et al. 2022). The limitations of BP have spurred research into more biologically plausible and hardware-efficient alternatives (Shinozaki 2021).

To circumvent these constraints, several non-BP training techniques have been proposed, including Predictive Coding (Rao and Ballard 1999; Salvatori et al. 2022), Evolutionary Algorithms (Salimans et al. 2017), non-BP Reinforcement Learning (RL) (Sutton and Barto 2018), and Competitive Hebbian Learning (Miconi 2021). While these strategies demonstrate better handling of black-box models and vanishing/exploding gradients, they are often computationally demanding, require substantial memory, and present other limitations. A recent development in this domain is the Forward-Forward algorithm (FF) (Hinton 2022), a non-BP technique involving two forward passes - one with positive (real) data, and the other with negative data. FF offers notable computational benefits in scenarios where the precise details of the forward computation are not known without the necessity of storing neural activities or interrupting for error derivative propagation. It is particularly advantageous for applications involving low-power analog hardware. Despite its advantages, the FF algorithm also exhibits slow convergence, inadequate performance on complex tasks, and the necessity of generating negative data.

This paper presents an approach to enhance the Forward-Forward Algorithm and overcome its shortcomings by concurrently extracting and separating features through Convolutional Neural Networks (CNN). Our framework results in improved representation learning and expedited training. Specifically, our main contributions are the following:

- Construct a layer-wise loss function that exploits channel-wise competitive learning. It reformulates the goodness function to eliminate the need for negative data construction and empowers each CNN layer to function as a standalone classifier.
- Introduce the Channel-wise Feature Separator and Extractor (CFSE) block that partitions the feature space via channel-wise grouped convolutional layers and simultaneously enables the network to learn compositional fea-

tures via standard non-separable convolutional layers.

Our approach is benchmarked against FF-based and non-Backpropagation methods reported in the literature on image classification tasks. Our method outperforms all FF-based models on basic image classification tasks, demonstrating quicker convergence rates and significantly better performance, achieving testing errors of 0.58%, 7.69%, 21.89%, and 48.77% on MNIST (LeCun et al. 1998), Fashion-MNIST (Xiao, Rasul, and Vollgraf 2017), CIFAR-10, and CIFAR-100 (Krizhevsky and Hinton 2009) respectively. Moreover, it significantly outperforms almost all the non-BP models, with SoftHebb (Journé et al. 2023) being the only competitive one.

Background and Related Work

The BP algorithm has seen numerous improvements in the literature. Techniques like batch normalization (Ioffe and Szegedy 2015), dropout, and weight initialization have been employed to mitigate issues such as overfitting, vanishing gradients, and initialization bias. Moreover, alternative optimization methods such as Adam (Kingma and Ba 2015) and Adagrad (Duchi, Hazan, and Singer 2011) demonstrate improved convergence and generalization. Methods that address symmetric forward and backward weights, waiting times, and reliance on expensive-to-generate labeled datasets are explored to tackle the biological implausibility and hardware inefficiency of BP (Sze et al. 2017). Self-supervised learning schemes, for instance, adjust network weights through self-generated supervisory errors propagated from the system’s output backward in an end-to-end manner (Chen et al. 2020), though these are constrained by the efficiency of distributed computing hardware.

In response to the limitations of traditional deep learning algorithms based on BP, various non-BP training schemes have been proposed in the literature. Predictive Coding operates on the principle of the brain’s predictive model, which continuously generates predictions and compares them with incoming sensory data (Song and Ermon 2020; Ororbia and Mali 2023). Similarly, Evolutionary Algorithms (Salimans et al. 2017) adopt the concept of natural selection, with the fittest individuals (or weights in the context of ANNs) selected for reproduction. Both techniques, unlike BP, are not end-to-end training processes, and are better in managing black-box models and the vanishing/exploding gradients problem compared to BP, despite having no perfect knowledge of forward pass computation (Barbalau et al. 2020). However, these methods are computationally demanding, memory-intensive, and, in the case of Predictive Coding, struggle with noise handling (Rao and Ballard 1999; Ofner, Millidge, and Stober 2022). A form of Reinforcement Learning can manage black-box models by correlating random perturbations to weights or neural activities with changes in a payoff function. However, its high variance and computational expense make it difficult to scale to larger networks (Hinton 2022; Sutton and Barto 2018).

Feedback Alignment (FA) (Lillicrap et al. 2016) offers a notable alternative to traditional backpropagation. Instead of fixed, symmetrical weights between forward and back-

ward passes, FA uses random initial feedback weights that the network learns to optimize. This enables independent training of hidden layers, starting from zero initial conditions. FA contributes to biologically plausible models by using mostly local error signals and eliminating the need for symmetrical weights. Its extension, Direct Feedback Alignment (Nøkland 2016; Frenkel, Lefebvre, and Bol 2021), further simplifies weight updates by avoiding non-local computations. Competitive Hebbian Learning enables neurons to compete for activation, enhancing the ones most responsive to input (Miconi 2021; Nessler, Pfeiffer, and Maass 2008). While effective in overcoming overfitting and task-specific constraints, it suffers from slow convergence and parameter tuning with SoftHebb (Journé et al. 2023) promising higher convergence rates and accuracy by training deep networks without feedback signals.

The FF algorithm (Hinton 2022) has also emerged as a compelling alternative to the traditional BP method. It uses two forward passes, one with positive (real) data and the other with negative data, i.e., fake data generated by the network. One approach for generating a positive-negative pair involves the creation of a vector that combines the correct labels and corresponding data for positive examples while generating “bad data” by pairing input images with incorrect labels. The primary objective of each network layer under this framework is to maximize the model’s performance on positive data and minimize it on negative data. This is achieved by optimizing a proposed “goodness” function, initially defined as the sum of squared activations within a layer. The layer-wise training structure of the FF algorithm enables the evaluation of each layer’s performance and facilitates a more transparent selection of hyperparameters and model architecture. Furthermore, the FF algorithm can operate without knowledge of forward computation details or the need for storing neural activities or halting error derivatives. The FF algorithm presents three main challenges. First, the need for negative data is highly task and dataset-specific. Second, the initial formulation of the goodness function is preliminary and calls for further refinement. Third, the FF algorithm is slower than BP and does not generalize as well on certain straightforward problems, which in practice renders BP more suitable for more complex models and datasets. Recent FF-based approaches have emerged that try to overcome the limitations of the original FF learning framework. The PFF-RNN model (Ororbia and Mali 2023) employs a dynamic FF recurrent neural system, integrating lateral competition, noise injection, and predictive coding to conduct credit assignments in neural systems. The Cascaded Forward (CaFo) model (Zhao et al. 2023), amends the original FF algorithm by introducing convolutional layer blocks and directly outputting label distributions for each cascaded block, thereby eliminating the need for negative data, while however showing extremely slow convergence rates.

For the context of FF learning, this work improves the state-of-the-art by (i) enhancing the network’s capacity to learn intraclass features and improve the accuracy and convergence rates of the FF algorithm, (ii) eliminating the necessity for negative data construction and broadening the applicability of layer-wise learning, and (iii) exploit the layer-

wise nature of FF learning to introduce Channel-wise Competitive Learning making it more applicable to CNNs, unlike competitive Hebbian learning methods that are restricted to competition among shared kernels.

Proposed Approach

This section outlines key components of our competitive channel-wise learning framework. These include the redefined *Convolutional Goodness*, and *PvN and CwC loss functions* for exploiting Channel-wise Competitive Learning, the model architecture featuring the introduced *CFSE block*, three types of *Predictors*, and the *Interleaved Layer Training strategy* governing the learning process.

Channel-wise Competitive Learning

Formulating Convolutional Goodness. We redefine *positive* and *negative* goodness to signify the activations of the channel groups associated with the target and the other classes, respectively, in order to evaluate the performance of convolutional layers. The feature map activation outputs of each convolutional layer are denoted as $\mathbf{Y} \in \mathbb{R}^{N \times C \times H \times W}$, where N is the number of samples in the mini-batch, C is the channel dimension, and H, W are the spatial dimensions. Assuming J data classes, \mathbf{Y} is partitioned into J subsets of $S = C/J$ channels, such that $\hat{\mathbf{Y}}_j \in \mathbb{R}^{N \times S \times H \times W}$ represents the j^{th} class. Then, the *holistic goodness factor* for each layer is a matrix, $\mathbf{G} \in \mathbb{R}^{N \times J}$. Each element of the matrix $G_{n,j}$ involves the computation of the goodness of each class j and example n , through the spatial and channel-wise mean of the square of $\hat{\mathbf{Y}}_j$ as

$$G_{n,j} = \frac{1}{S \times H \times W} \sum_{s=1}^S \sum_{h=1}^H \sum_{w=1}^W \hat{Y}_{n,j,s,h,w}^2. \quad (1)$$

We define the binary label mask for a mini-batch of N examples with J classes that can be represented as a matrix $\mathbf{Z} \in \{0, 1\}^{N \times J}$, where each row corresponds to an example and each column corresponds to a class. Letting $z_n \in \{1, \dots, J\}^N$, denote the target class of example n , the elements of \mathbf{Z} are defined as

$$\text{Binary Label Mask } Z_{n,j} = \begin{cases} 1, & \text{if } z_n = j, \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

The *positive goodness*, $\mathbf{g}^+ \in \mathbb{R}^N$, is a vector that represents the goodness of the subset of channels associated with the target class $j = z_n$. It is calculated through the dot product of \mathbf{G} and the transpose of the binary label mask, \mathbf{Z}^T as

$$\mathbf{g}^+ = \mathbf{G} \cdot \mathbf{Z}^T. \quad (3)$$

The *negative goodness*, $\mathbf{g}^- \in \mathbb{R}^N$, is a vector that denotes the goodness values of the remaining channel blocks for the non-target categories, where $j \neq z_n$ calculated as

$$\mathbf{g}^- = \mathbf{G} \cdot (\mathbf{1} - \mathbf{Z}^T). \quad (4)$$

PvN Loss Function. We initially investigated the use of redefined positive and negative goodness, \mathbf{g}^+ and \mathbf{g}^- , in a sigmoidal, cross-entropy loss function, L_{PvN} . This function is built on the FF positive and negative framework and is designed to maximize \mathbf{g}^+ and minimize the mean negative goodness obtained by dividing \mathbf{g}^- with the number of non-target categories, $J-1$. Similar to the original FF, we employ a hyperparameter, θ , as a threshold determining the boundary between the correct and the false class, formulated as

$$L_{PvN} = \frac{1}{2N} \sum_{n=1}^N \left[\log(1 + \exp(-g_n^+ + \theta)) + \log\left(1 + \exp\left(\frac{1}{(J-1)}g_n^- - \theta\right)\right) \right] \quad (5)$$

CwC Loss Function. We then move on to establish a channel-wise loss function, L_{CwC} , to promote competitive learning through the channel dimension. It does this using the different goodness scores, calculated for each class, as logits in a Softmax-based cross-entropy loss. This unique approach enables each convolutional layer to function as an independent classifier, facilitating layer-wise performance evaluation and yielding a more transparent learning process compared to traditional 'black box' methods. The Channel-wise Competition loss function, L_{CwC} , trains each convolution layer as a classifier by redefining the Softmax function to operate with goodness scores. It does this using the positive goodness, \mathbf{g}_n^+ , for each sample n in the batch to estimate the logit that represents the probability of the target class over the total goodness score for each sample. Normalizing the goodness scores assists in the creation of a competitive dynamic among classes encouraging the model to increase confidence in the correct prediction through \mathbf{g}_n^+ , while simultaneously suppressing the scores of incorrect classes.

$$L_{CwC} = -\frac{1}{N} \sum_{n=1}^N \log\left(\frac{\exp(g_n^+)}{\sum_{j=1}^J \exp(G_{n,j})}\right) \quad (6)$$

Network Architecture

CFSE block. The main ideas presented in the previous section are instantiated through a dedicated architecture that leverages channel-wise grouped convolutional layers to partition the feature space along the channel dimension. This not only improves efficiency and performance but also allows for improved representation learning by enabling the vectors formed for each class to learn compositional features from different kernels. The network architecture comprises a series of Convolutional Feature Space Separation and Extraction (CFSE) blocks, each containing two convolutional layers for feature extraction and space separation, as depicted in Figure 1. The first layer is a standard non-separable convolution (Conv2D) layer, which is used to learn joint depth features of the input data and joint inter-class features of different subsets of channels in successive CFSE blocks. The second layer is a channel-wise grouped convolution (GroupConv) layer, which enhances the extraction of intra-class features and increases the gap between the different subsets of features. Each layer is trained independently in a layer-wise manner. As illustrated in Figure 1, the

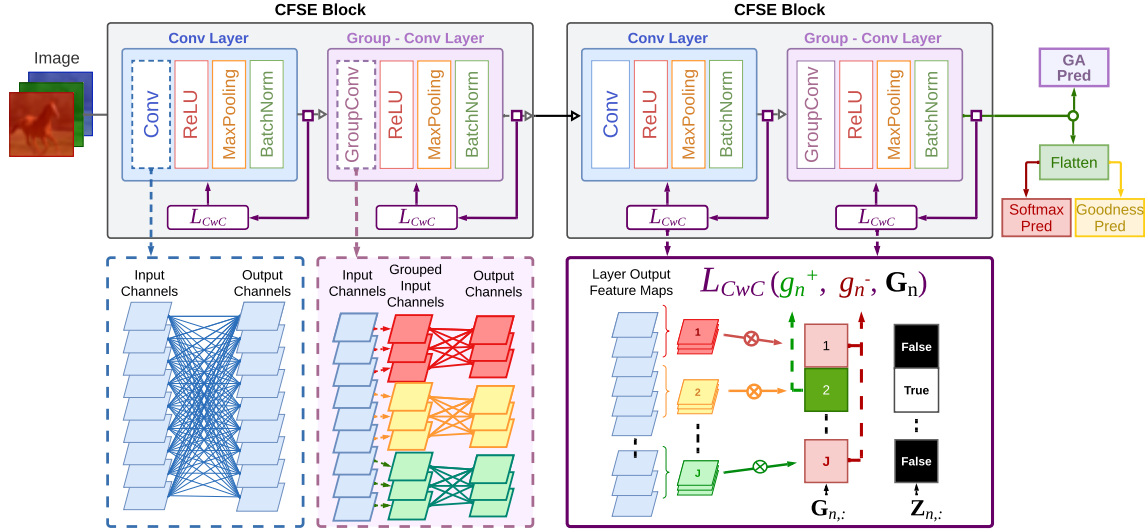


Figure 1: Top: Network Architecture, Components of CFSE Block, Standard Convolutional (Conv) Layer, Grouped Convolutional (Group - Conv) Layer; Bottom: Non-separable Convolution, Channel-wise Group - Conv, CwC Loss Function.

output features of the convolutional operations are passed through ReLU activations and batch normalization (Batch-Norm) (Ioffe and Szegedy 2015), with max-pooling applied to the ReLU activations of layers operated on GroupConv.

Channel-wise Grouped Convolutions. The CFSE blocks employ channel-wise grouped convolutions on the feature maps to achieve feature space separation through the channel dimension. We utilize Grouped Convolutions to allow for intra-class specific feature learning, by enabling each subset of the input and output channels to represent a unique class. This helps the network to learn more specialized features for each class, thereby enhancing the overall accuracy and classification performance. The implementation of grouped convolutions contributes to a decrease in the requisite number of parameters and computations for each layer, thereby enhancing efficiency and accelerating the training process (Guo et al. 2020), (Schuler et al. 2021). This is substantiated through the comparative analysis of primary architectures utilized in this study, presented in Table 2.

Following the notations used in (Guo et al. 2020), we denote the weights of a GroupConv layer, \mathbf{W}_{GC} , as a tensor of size $C_{out} \times k \times k \times C_{in}$, where C_{out} and C_{in} are the number of output channels and input channels, respectively, and k is the kernel size. Splitting of \mathbf{W}_{GC} through the channel dimension into J subsets, each associated with a specific class denoted as $\hat{\mathbf{S}}_1, \hat{\mathbf{S}}_2, \dots, \hat{\mathbf{S}}_J$, so that each subset, $\hat{\mathbf{S}}_j$, is comprised of C_{out}/J and C_{in}/J channels, where j denotes the j^{th} class. Let $\hat{\mathbf{S}} = \text{diag}(\hat{\mathbf{S}}_1, \hat{\mathbf{S}}_2, \dots, \hat{\mathbf{S}}_J)$, such that it is a block-diagonal matrix, the channel-wise grouped convolution operation can be formulated as

$$\hat{y}_j = \hat{\mathbf{S}}_j x_j + b_j, j = 1, \dots, J, \quad (7)$$

where x_j and \hat{y}_j are the input and output tensors of the j^{th} subset of sizes $k \times k \times C_{in}/J$ and C_{out}/J respectively. It is important to note that Grouped Convolutions require the

C_{out} and C_{in} to be divisible by the number of subsets, J , and each subset has the same number of channels.

Predictors. We utilize three types of predictors: (i) the FF-Goodness predictor (Goodness-pred), (ii) the Softmax predictor (Softmax-pred), and (iii) the Global Averaging predictor (GA-pred). The first two predictors operate on the flattened output of the last CFSE block in order to classify the input sample. The Softmax predictor is a standard Softmax layer trained with Cross-Entropy loss using the unnormalized outputs of the last CFSE Block. For the evaluation process, we select the class with the highest confidence as our prediction. The Goodness Predictor is a series of 2 dense layers with 1024 neurons each. It follows the paradigm of (Hinton 2022) where the target label is overlaid as an embedded annotation on the flattened feature maps as positive features, and the non-target label as negative features while using the mean of the normalized squared activations as goodness. Similar to the FF algorithm, we normalize the length of the hidden vector before passing it to the next layer to force it to rely on the relative activities of the neurons in the previous layer and not their magnitude. For the evaluation process, we utilize the original inference methodology in (Hinton 2022), and run the dense layers for each label individually, aggregating the goodness values across all layers. The label associated with the highest cumulative goodness is then selected. GA-pred negates the need for Fully Connected (FC) layers and computes the mean squared value for each class subset, selecting the class with the highest mean as the prediction. This renders GA-pred method more lightweight and flexible, given that FC layers usually contribute to the largest number of parameters as shown in Table 2.

Interleaved Layer Training (ILT) Strategy

Capitalizing on the modular structure of our network, we seek to facilitate better and quicker convergence. ILT, en-

Algorithm 1: Interleaved Layer Training (ILT) Strategy

Require: Arrays $startEp$; $plateauEp$ of size $layerCount$

- 1: Initialize $startEp[:] = 0$; $plateauEp[:] = maxEpoch$
- 2: % Identify Optimal Start and Training Plateau for each Layer
- 3: **for** $i = 1$ to $layerCount$ **do**
- 4: **for** $epoch = 1$ to $maxEpoch$ **do**
- 5: **for** $j = 1$ to i **do**
- 6: **if** $startEp[j] \leq epoch \leq plateauEp[j]$ **then**
- 7: Train j^{th} layer until it reaches a plateau
- 8: $plateauEp[i] \leftarrow epoch$
- 9: **if** $fastMode$ **then**
- 10: $startEp[i + 1] \leftarrow epoch - N$
- 11: % Perform interleaved training
- 12: **for** $epoch = 1$ to $max(plateauEp)$ **do**
- 13: **for** $i = 1$ to $layerCount$ **do**
- 14: **if** $startEp[i] \leq epoch \leq plateauEp[i]$ **then**
- 15: Train i^{th} layer

tails the timed initiation and termination of training epochs for each layer, allowing parallel training with preceding layers for a certain number of epochs while also allowing for independent layer training for a different number of epochs. We determined the appropriate ILT for each layer through the process outlined in Algorithm 1. We start from the first layer and train until it reaches a plateau which we use as the stopping epoch. For each successive layer, we restart the training and record their plateaus and stopping epochs as well. Then we performed interleaved training where all layers are trained together. However, instead of training all the layers continuously, we select to stop each layer at the noted plateau epoch. This strategy reduces the chances of stagnation in local minima while allowing each layer to be fine-tuned on constant feature outputs from the previous layer. Regarding the choice of when to start training each layer, we identified through trial and error that starting the training for all layers simultaneously yields the best accuracy results. However, the fastest convergence was achieved when we started training each successive layer, a number of epochs, N , before the plateau of the predecessor. Moreover, N differed and tended to be larger for more complex datasets. While this learning strategy provides better results it is still preliminary and there is substantial room for further refinement and enhancement.

Experimental Design and Results

This section details the experiments carried out to evaluate the proposed techniques and compare them with existing works. Specifically, the performance of the various models was evaluated on three benchmark datasets commonly used in literature, namely MNIST (LeCun et al. 1998), Fashion-MNIST (Xiao, Rasul, and Vollgraf 2017), and CIFAR-10 (Krizhevsky and Hinton 2009). Our models were trained on an NVIDIA-RTX 4080 GPU and the i9-13900K Intel CPU for 20 epochs on MNIST, 50 epochs on Fashion-MNIST, and CIFAR-10. In the experiments, the learning rate used for both the convolutional and dense layers was set to 0.01, and the batch size to 128. For comprehensive details on the model configurations, refer to the supplementary material.

Ablation Study

Our ablation study is designed to demonstrate the impact of various model components, including our proposed loss functions, architectural modifications, and learning frameworks. The structure and performance of each model, illustrated in Table 1, provide detailed insights into their individual contributions. We first replicate the foundational FF Fully Connected (FC) model, FF-FC(rep*), as outlined in (Hinton 2022) and use the original negative data construction method. The FF-FC_Hybrid uses the hybrid negative data generation process which involves the fusion of two distinct images through a large-area binary mask, which is generated from a blurred and thresholded random bit image. Subsequent models incorporate CNNs, beginning with FF-CNN_Hybrid which employs the same negative data generation strategy and goodness function as FF-FC_NegData, with each feature map element serving as an individual neuron. Further configurations differ by the type of loss function used, CwC and PvN, to train the Convolutional Layers, and the predictors - Goodness Predictor (Gd), Softmax Predictor (Sf), and Global Averaging Predictor (GA). Finally, the CFSE models explore the configurations that use the CFSE architecture. The naming convention contains first the architecture (FF-FC/FF-CNN/CFSE), the loss function (PvN/CwC), and the predictor (Sf/Gd/GA).

Impact of Loss Function. The choice of loss function influences the learning dynamics and subsequently the performance of the models. Our ablation study reiterates this fact, as seen through the performance variance between models that utilized the CwC and PvN loss functions. We delve deeper into this analysis by comparing the performance of three models - FF-CNN_Hybrid, FF-CNN_PvN+Gd, and FF-CNN_CwC+Gd - that differ primarily in the choice of loss function they employ. The model FF-CNN_Hybrid, which uses the original goodness function, demonstrates an average performance with test errors of 2.27%, and 19.21% on the MNIST, and Fashion-MNIST datasets respectively. The FF-CNN_PvN+Gd model employs the PvN loss function with a notable improvement in test errors on Fashion-MNIST 12.73%. The most dramatic improvement is obtained by switching to the CwC loss function with the testing errors reduced to 0.91%, 7.88%, and 23.16% on MNIST, Fashion-MNIST, and CIFAR-10, respectively. This comparative analysis highlights the enhanced performance brought about by the CwC and PvN loss functions, with the CwC leading to the most substantial improvements across all datasets and model configurations. Moreover, it is intriguing to note that with the increase in the complexity of the dataset from MNIST to CIFAR-10, the impact of the loss function on model performance becomes more pronounced demonstrating CwC superiority for more complex tasks.

Impact of CFSE. The adoption of the CFSE architecture leads to complexity and efficiency improvements. By comparing FF-CNN with the CFSE configurations, we notice slight reductions in test errors and training epochs on the most complex dataset, CIFAR-10. The computational benefits are prevalent in Table 2, where we present the total parameters and the number of multiplication-addition opera-

Model	Negative Data	PvN Loss	CwC Loss	CNN Layers	CFSE Blocks	Predictor	Testing Error (%)		
							MNIST	Fashion-MNIST	CIFAR-10
FF-FC(rep*)	✓					Gd	2.02 ± 0.37	14.19 ± 1.91	46.97 ± 1.32
FF-FC_Hybrid	✓					Gd	2.06 ± 0.31	19.72 ± 0.28	-
FF-CNN_Hybrid	✓			✓		Gd	2.27 ± 0.37	19.21 ± 1.50	-
FF-CNN_PvN+Gd		✓		✓		Gd	2.34 ± 0.04	12.73 ± 0.46	33.79 ± 0.84
FF-CNN_PvN+Sf		✓		✓		Sf	2.38 ± 0.02	12.42 ± 0.13	32.98 ± 0.16
FF-CNN_CwC+GA			✓	✓		GA	0.86 ± 0.04	8.79 ± 0.17	25.01 ± 0.24
FF-CNN_CwC+Gd			✓	✓		Gd	0.91 ± 0.16	7.88 ± 0.53	23.16 ± 0.16
FF-CNN_CwC+Sf			✓	✓		Sf	0.59 ± 0.09	7.77 ± 0.53	22.54 ± 0.16
CFSE_CwC+GA			✓	✓	✓	GA	0.86 ± 0.13	8.91 ± 0.01	27.25 ± 0.49
CFSE_CwC+Gd			✓	✓	✓	Gd	0.59 ± 0.02	7.95 ± 0.12	24.43 ± 0.34
CFSE_CwC+Sf			✓	✓	✓	Sf	0.58 ± 0.08	7.69 ± 0.32	21.89 ± 0.44

Table 1: Detailed Configuration of Different Models and their Performance used in the ablation study. The models vary by the loss function, CNN architecture, and the type of predictor used. The performance comparison is in terms of test error percentage.

tions for each architecture equipped with an Sf or GA Predictor, using CIFAR-10 image dimensions of [32, 32, 3] as input. The parameters and operations were calculated with the open-source toolbox, torchinfo (Harris 2016). The CFSE architecture demonstrates lower computational complexity compared to both FFCNN models, necessitating a fraction of its operations. This implies that given equivalent conditions, CFSE would be faster and more resource-efficient for training and execution. The deployment of grouped convolutions within the CFSE architecture is a primary contributor to this enhanced efficiency, reducing the computational cost.

Qualitative CFSE-CwC Results. We examine feature maps generated by the CFSE-CwC’s CNN layers on various MNIST test inputs. Feature maps for each subset are computed as the mean activations over the channel-dimension. As shown in rows 3 and 4 of Figure 2, the activations that correspond to the true class for both examples exhibit higher values proving that the channel-wise competition promoted by L_{CwC} enables the CNN layers to separate the feature space and act as independent classifiers. Moreover, following the output from the initial CFSE block (Group-Conv Layer), a clearer separation of the feature space is noticeable, underscored by a larger disparity between the activations of the target and the non-target class. This verifies that GroupConv layers enhance the extraction of intra-class features as is also shown by the training plots in the supplementary material.

Predictor Performance. The choice of predictor, hence the way that the final classification is computed significantly impacts model performance since each predictor demonstrates distinct performance characteristics. As shown in Table 3, the models that use an Sf predictor consistently outperform the models that use the other two predictors across all datasets, with Gd performing on average better than GA. Sf and Gd perform better but require a dedicated FC layer which not only increases the number of parameters by a factor of two as shown in 2, but also worsens the convergence rate. There is a clear trade-off with respect to accuracy and computational expense between the FC Predictors and GA.

Comparison with Related Work

Table 3 illustrates the results of our experiments, where our top performing model, CFSE.CwC+Sf is compared with recent related work approaches. Our model’s architecture on all datasets, excluding CIFAR-100, comprised two CFSE Blocks and is configured as in Table 2. We allowed longer training time for CIFAR-100 and increased the complexity to 3 CFSE Blocks. Additionally, we ventured into Class-grouping, addressing the scalability constraint, training the first two CFSE blocks on the 20 super-classes, in line with their designation in the official CIFAR dataset, and the third block across all 100 classes. For both configurations Max pooling was used only on the Group-Conv layer ReLU outputs. We utilized the ILT Learning strategy with simultaneous initialization for all layers. We analyze its comparative advantages within three major learning paradigms: FF-Based Methods, Non-BP Methods, and BP methods.

The FF-based models include the original Forward-Forward (FF-FC*) model, the Cascaded Forward algorithm (CaFo) Zhao et al. (2023), and Predictive FF with Recurrent NN (PFF-RNN). Our model consistently outperforms all FF-based methods across all datasets. On the MNIST dataset, our model reports less than half the testing error of its nearest FF competitors, and on CIFAR-10 half the error of the original FF-FC. The difference in performance with FF-FC is even greater on CIFAR-100 demonstrating that our approach advances FF to handle more complex datasets effectively. Our model significantly outperforms the second-best, CaFo, despite CaFo requiring 5000 epochs for training.

Remarkably, our method also excels in comparison to mature non-Backpropagation learning frameworks, namely, the FFF algorithm (Flügel et al. 2023), Direct Feedback Alignment (DFA) (Nø kland 2016), DRTP (Frenkel, Lefebvre, and Bol 2021) which show higher testing error and slower convergence. The only competitive method is SoftHebb (Journé et al. 2023) which employs advanced techniques like adaptive learning rates for a performance boost. Despite the absence of such mechanisms, our model attains comparable or better performance on all benchmarked datasets.

Finally, we compare our model with the BP model implemented in (Hinton 2022) and a Backpropagation Convolutional Model, BP-CNN(rep*), which shares the same

Model	CNN Layers	Group Conv	Kernel (size, stride)	Maxpool (Kernel 2x2, Stride 2)	Total Parameters	Total mult-adds (M)
CFSE_CwC+Sf	4 (2 Blocks)	[no, yes, no, yes]	3, 1	[no, yes, no, yes]	588,133	73.4
CFSE_CwC_GA	4 (2 Blocks)	[no, yes, no, yes]	3, 1	[no, yes, no, yes]	280,920	73.1
FF-CNN_CwC+Sf	4	[no, no, no, no]	3, 1	[no, yes, no, yes]	1,534,210	325.6
FF-CNN_CwC_GA	4	[no, no, no, no]	3, 1	[no, yes, no, yes]	1,227,000	325.2

Table 2: Comparison between the different architectures on the Number of Parameters and Multiplication-Addition operations with [32, 32, 3] image dimensions as input and layer channels of [20, 80, 240, 480]




Figure 2: Feature Map Visualization generated by the CFSE_CwC's CNN layers on MNIST inputs (Left:7; Right:4). Rows correspond to layers (Block1: Conv, GroupConv; Block2: Conv, GroupConv), and columns to the subsets of channels tied to each class. The activations that correspond to the true class for both examples exhibit higher values on the last two layers.

Model	MNIST		Fashion-MNIST		CIFAR-10		CIFAR-100	
	Test Er. (%)	Epochs	Test Er. (%)	Epochs	Test Er. (%)	Epochs	Test Er. (%)	Epochs
BP-FC (*)	1.4	40	NR	-	39	40	NR	-
BP-CNN (rep*)	0.56	20	6.75	30	13.63	50	43.38	100
FFF	2.84	100	NR	-	53.96	200	NR	-
DRTP	1.5	100	NR	-	31.04	200	NR	-
DFA	1.02	300	NR	-	26.9	300	59.0	300
SoftHebb	0.65	50	NR	-	19.7	100	44.0	200
FF-FC (*)	1.37	60	NR	-	41	-	NR	-
PFF-RNN	1.3	60	10.41	100	NR	-	NR	-
CaFo	1.3	5000	NR	-	32.57	5000	59.24	5000
CFSE_CwC+Sf (ours)	0.58	10 (20)	7.69	25 (50)	21.89	40 (50)	48.77	150 (700)

NR* - No Results on the specific benchmarks

Table 3: Performance Comparison (Testing Error (%) and Training Time (Epochs)) of Proposed Method with recent non-backpropagation and FF-inspired approaches found in literature, as well as Backpropagation approaches that follow our architecture across different datasets. The epochs in parenthesis represent the training time needed for the Sf-Predictor.

architecture as CFSE_CwC+Sf. Though BP-CNN exhibits slightly better performance on average, our model narrows the gap significantly achieving comparable performance on MNIST, Fashion-MNIST, and CIFAR-100. This indicates that our approach presents a compelling alternative to traditional BP methods. These findings underscore the effectiveness and simplicity of our approach, which not only excels in testing accuracy but also in convergence rate. The results indicate substantial opportunities to further explore CNN design and loss functions within the FF learning paradigm.

Conclusion

This work has presented a novel approach to address the limitations of the Forward-Forward algorithm. A loss function inducing competitive learning between class-specific features has been introduced along with a block and network architecture that utilize grouped convolutions for reduced

parameters and efficiency. Experimental results on four image classification datasets exhibit significantly higher performance compared to existing FF methods, indicating a promising direction for further research.

Limitations. We present a competitive learning approach and architecture for improving various aspects of the recently introduced FF algorithm demonstrating improved performance over existing FF methods. Nevertheless, there are some potential limitations that warrant further investigation. The approach still lacks behind backpropagation in more complex datasets. Another potential limitation is scalability to a larger number of classes when dealing with a varying dataset. We assume that it would be possible to have a more elaborate network architecture and features shared between similar classes with adaptive class grouping to alleviate this.

Acknowledgments

This work is supported by the European Union (i. Horizon 2020 Teaming, KIOS CoE, No. 739551, and ii. ERC, URANUS, No. 101088124) and from the Government of the Republic of Cyprus through the Deputy Ministry of Research, Innovation, and Digital Policy.

Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

- Barbalau, A.; Cosma, A.; Ionescu, R. T.; and Popescu, M. 2020. Black-Box Ripper: Copying Black-Box Models Using Generative Evolutionary Algorithms. In *Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS 2020)*, 20120–20129. online: Curran Associates, Inc.
- Caruana, R.; Lawrence, S.; and Giles, C. 2000. Overfitting in Neural Nets: Backpropagation, Conjugate Gradient, and Early Stopping. In Leen, T.; Dietterich, T.; and Tresp, V., eds., *Proceedings of the 13th International Conference on Neural Information Processing Systems, (NIPS 2000)*. Denver, CO: MIT Press.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. E. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In *Proceedings of the 37th International Conference on Machine Learning, (ICML 2020)*, volume 119, 1597–1607. Vienna, Virtual Event: PMLR.
- Duchi, J.; Hazan, E.; and Singer, Y. 2011. Adaptive Sub-gradient Methods for Online Learning and Stochastic Optimization. *Journal of Machine Learning Research*, 12(61): 2121–2159.
- Flügel, K.; Coquelin, D.; Weiel, M.; Debus, C.; Streit, A.; and Götz, M. 2023. Feed-Forward Optimization With Delayed Feedback for Neural Networks. arXiv:2304.13372.
- Frenkel, C.; Lefebvre, M.; and Bol, D. 2021. Learning Without Feedback: Fixed Random Learning Signals Allow for Feedforward Training of Deep Neural Networks. *Frontiers in Neuroscience*, 15: ISSN: 1662–453X.
- Guo, Q.; Wu, X.-J.; Kittler, J.; and Feng, Z. 2020. Self-grouping Convolutional Neural Networks. *Neural Networks*, 132: 491–505.
- Harris, T. 2016. torchsummary: A PyTorch tool for network summary and visualization. <https://github.com/sksq96/pytorch-summary>. Accessed: 2023-05-16.
- Hinton, G. E. 2022. The Forward-Forward Algorithm: Some Preliminary Investigations. arXiv:2212.13345.
- Ioffe, S.; and Szegedy, C. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning, (ICML 2015)*, 448–456. Lille: PMLR.
- Journé, A.; Rodriguez, H. G.; Guo, Q.; and Moraitis, T. 2023. Hebbian Deep Learning Without Feedback. In *Proceedings of the Eleventh International Conference on Learning Representations (ICLR 2023)*. Kigali: OpenReview.net.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In Bengio, Y.; and LeCun, Y., eds., *Proceedings of the 3rd International Conference on Learning Representations, (ICLR 2015)*. San Diego, CA: OpenReview.net.
- Krizhevsky, A.; and Hinton, G. 2009. Learning Multiple Layers of Features from Tiny Images. Technical Report 0, University of Toronto, Toronto, Ontario.
- Lawrence, S.; and Giles, C. 2000. Overfitting and neural networks: conjugate gradient and backpropagation. In *Proceedings of the 2000 IEEE-INNS-ENNS International Joint Conference on Neural Networks (IJCNN 2000)*, volume 1, 114–119.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-Based Learning Applied to Document Recognition. *Proceedings of the IEEE*, 86(11): 2278–2324.
- Lillicrap, T. P.; Cownden, D.; Tweed, D. B.; and Akerman, C. J. 2016. Random synaptic feedback weights support error backpropagation for deep learning. *Nature communications*, 7(1): 13276.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *Proceedings of the 6th International Conference on Learning Representations, (ICLR 2018)*. Vancouver: OpenReview.net.
- Miconi, T. 2021. Hebbian learning with gradients: Hebbian convolutional neural networks with modern deep learning frameworks. arXiv:2107.01729.
- Moraitis, T.; Toichkin, D.; Journé, A.; Chua, Y.; and Guo, Q. 2022. SoftHebb: Bayesian Inference in Unsupervised Hebbian Soft Winner-Take-All Networks. *Neuromorphic Computing and Engineering*, (4): 044017.
- Nessler, B.; Pfeiffer, M.; and Maass, W. 2008. Hebbian Learning of Bayes Optimal Decisions. In *Proceedings of the 21st International Conference on Neural Information Processing Systems, (NIPS 2008)*. Curran Associates, Inc.
- Nø kland, A. 2016. Direct Feedback Alignment Provides Learning in Deep Neural Networks. In Lee, D.; Sugiyama, M.; Luxburg, U.; Guyon, I.; and Garnett, R., eds., *Proceedings of the 29th International Conference on Neural Information Processing Systems, (NIPS 2016)*, volume 29. Curran Associates, Inc.
- Ofner, A.; Millidge, B.; and Stober, S. 2022. Generalized Predictive Coding: Bayesian Inference in Static and Dynamic Models. In *Shared Visual Representations in Human and Machine Intelligence Workshop (SVRHM 2022) @ NeurIPS*.
- Ororbia, A.; and Mali, A. 2023. The Predictive Forward-Forward Algorithm. arXiv:2301.01452.
- Rao, R. P.; and Ballard, D. H. 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, 2(1): 79–87.

- Salimans, T.; Ho, J.; Chen, X.; and Sutskever, I. 2017. Evolution Strategies as a Scalable Alternative to Reinforcement Learning. arXiv:1703.03864.
- Salvatori, T.; Song, Y.; Xu, Z.; Lukasiewicz, T.; and Bogacz, R. 2022. Reverse differentiation via predictive coding. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI 2022)*, 7, 8150–8158. Virtual Event.
- Schuler, J. P. S.; Romani, S.; Abdel-Nasser, M.; Rashwan, H. A.; and Puig, D. 2021. Grouped Pointwise Convolutions Significantly Reduces Parameters in EfficientNet. In *Proceedings of the 23rd International Conference of the Catalan Association for Artificial Intelligence, (CCIA 2021)*. Virtual Event.
- Shinozaki, T. 2021. Biologically motivated learning method for deep neural networks using hierarchical competitive learning. *Neural Networks*, 144: 271–278.
- Song, J.; and Ermon, S. 2020. Multi-label Contrastive Predictive Coding. In *Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS 2020)*, 8161–8173. online: Curran Associates, Inc.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement Learning: An Introduction*. Cambridge, MA: The MIT Press, second edition.
- Sze, V.; hsin Chen, Y.; Yang, T.-J.; and Emer, J. S. 2017. Efficient Processing of Deep Neural Networks: A Tutorial and Survey. *Proceedings of the IEEE*, 105: 2295–2329.
- Tan, Z.-H.; Xie, Y.; Jiang, Y.; and Zhou, Z.-H. 2022. Real-Valued Backpropagation is Unsuitable for Complex-Valued Neural Networks. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS 2022)*, 34052–34063. New Orleans, LA: Curran Associates, Inc.
- Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms. arXiv:1708.07747.
- Zhao, G.; Wang, T.; Li, Y.; Jin, Y.; Lang, C.; and Ling, H. 2023. The Cascaded Forward Algorithm for Neural Network Training. arXiv:2303.09728.