

©Copyright 2021

Jinlin Xiang

# Knowledge Distillation Circumvents Nonlinearity for Optical Convolutional Neural Networks

Jinlin Xiang

A thesis  
submitted in partial fulfillment of the  
requirements for the degree of

Master of Science

University of Washington

2021

Committee:

Eli Shlizerman

Xu Chen

Duane W. Storti

Program Authorized to Offer Degree:  
Mechanical Engineering

University of Washington

**Abstract**

Knowledge Distillation Circumvents Nonlinearity for Optical Convolutional Neural Networks

Jinlin Xiang

Chair of the Supervisory Committee:

Eli Shlizerman

Department of Electrical & Computer Engineering, Department of Applied Mathematics

In recent years, Convolutional Neural Networks (CNNs) have enabled ubiquitous image processing applications. As such, CNNs require fast runtime (forward propagation) to process high-resolution visual streams in real time. This is still a challenging task even with state-of-the-art graphics and tensor processing units. The bottleneck in computational efficiency primarily occurs in the convolutional layers. Performing operations in the Fourier domain is a promising way to accelerate forward propagation since it transforms convolutions into elementwise multiplications, which are considerably faster to compute for large inputs and kernels. Furthermore, such computation could be implemented using an optical  $4f$  system with orders of magnitude faster operation. However, a major challenge in using this spectral approach, as well as in an optical implementation of CNNs, is the inclusion of a nonlinearity between each convolutional layer, without which CNN performance drops dramatically. Here, we propose a Spectral CNN Linear Counterpart (SCLC) network architecture and develop a Knowledge Distillation (KD) approach to circumvent the need for a nonlinearity and successfully train such networks. While the KD approach is known in machine learning as an effective process for network pruning, we adapt the approach to transfer the knowledge from a nonlinear network (*teacher*) to a linear counterpart (*student*). We show that the KD approach can achieve performance that easily surpasses the standard linear version of a CNN and could approach the performance of the nonlinear network. Our simulations show that the possibility of increasing the resolution of the input image allows our proposed  $4f$  optical linear network

to perform more efficiently than a nonlinear network with the same accuracy on two fundamental image processing tasks: (i) object classification and (ii) semantic segmentation.

## TABLE OF CONTENTS

	Page
List of Figures . . . . .	ii
List of Tables . . . . .	iv
Chapter 1: Introduction . . . . .	1
Chapter 2: Methodology . . . . .	5
2.1 Knowledge Distillation (KD) . . . . .	5
2.2 The Spectral CNN Linear Counterpart (SCLC) . . . . .	7
2.3 Knowledge Distillation in SCLC . . . . .	12
2.4 Optical Implementation of SCLC . . . . .	13
Chapter 3: Computational Experiments . . . . .	15
3.1 Spectral Counterpart Layer Efficiency . . . . .	15
3.2 Classification . . . . .	16
3.3 Object Segmentation . . . . .	19
Chapter 4: Conclusion . . . . .	22
Bibliography . . . . .	24
Appendix A: Appendix Results . . . . .	29
A.1 SCLC Model . . . . .	29
A.2 STL Task . . . . .	31
A.3 Object Segmentation Task . . . . .	32

## LIST OF FIGURES

Figure Number	Page	
2.1	Illustration of KD training. Student network (green-bottom) parameters $\Phi$ are updated according to an interpolation of two losses: the Student loss (Crossentropy loss between data and the student model output) and Temp loss (Cross-entropy loss or KL divergence) between the teacher network (blue-top) and student class distribution with a temperature parameter $T$ . . . . .	6
2.2	Nonlinear CNN (top) and the proposed substitute, Spectral CNN Linear Counterpart (SCLC) (bottom). Top: An exemplar nonlinear CNN with layers that include operations of Convolution, Nonlinear RELU activation, and Max-Pool. Bottom: SCLC of the CNN shown in top row. The Convolution corresponds to the elementwise product in the spectral domain. The RELU operation is excluded. The Max-Pool layer is represented by a center-crop operation in the spectral domain. . .	8
2.3	Pooling in SCLC: Comparison for different pooling schemes in different dimension reductions. Top: Max pooling; Bottom: Spectral Pooling. The Spectral pooling retains more information on any arbitrary output dimension. . . . .	12
2.4	Optical $4f$ structure. The OSCLC can be implemented in $4f$ structure parallelly. .	14
3.1	Comparison of forward propagation time between AlexNet (blue line) and its SCLC (red line). Left: Propagation time through a single convolutional layer for variable image dimensions; Right: Propagation time through a single pooling layer for variable image dimensions. The dashed lines mark the resolution of an HD image and its corresponding runtime in AlexNet and SCLC. . . . .	16
3.2	Accuracy and Forward propagation time for classification task on High-10 dataset. Left: AlexNet and SCLC accuracy for varying input resolution; Right: Forward propagation runtime of AlexNet, SCLC, and OSCLC for varying input resolution.	18
3.3	Examples of object segmentation for the considered benchmarks of (i) Car Segmentation, (ii) Face Recognition, and (iii) VOC2012. Left to right: Input image from datasets, Ground truth, Our work (SCLC/ OSCLC). . . . .	21
A.1	Examples of object segmentation for the considered benchmarks of Car Segmentation. Left to right: Input image from datasets, Groundtruth, Our work (SCLC/ OSCLC) . . . . .	33

A.2 Examples of object segmentation for the considered benchmarks of VOC2012. Top to bottom: Input image from datasets, Groundtruth, Our work (SCLC/ OSCLC) . . . 33

## LIST OF TABLES

Table Number	Page
2.1 Comparison of forward propagation computational complexity between Nonlinear CNN, SCLC implemented on electronic hardware, SCLC implemented on optical hardware . . . . .	11
3.1 Accuracy and estimate the efficiency in forward propagation in classification . . . .	19
3.2 Ablation studies of pooling components, SCLC, backend and KD training . . . . .	20
3.3 Accuracy and forward propagation rate (frames per second) for object segmentation task tested on three benchmarks. . . . .	21
A.1 Implementation details of CNN, Optical fronted CNN and SCLC . . . . .	29
A.2 Implementation details of UNet (U shape CNN) and SCLC . . . . .	30
A.3 Accuracy for object classification task tested on STL-10 and Cifar-10 . . . . .	31
A.4 Subclass accuracy on Cifar-10 and STL-10 . . . . .	32

## **ACKNOWLEDGMENTS**

I wish to express sincere appreciation to the University of Washington, where I have had the opportunity to work with professor Eli Shlizerman and professor Arka Majumdar. This thesis would not be finished without help from them. They provided great guidance and feedback from last year. I also would thank professor Xu Chen and professor Duane W. Storti to help me formalize the final version of the thesis. Finally, great thanks to Washington Research Foundation, UW Reality Lab, Facebook, Google, Amazon, and Futurewei to support our work.

## **DEDICATION**

I dedicate this thesis to my parents and friends for their love and support.

## Chapter 1

### INTRODUCTION

Convolutional Neural Network (CNN) architectures are well known for their ability to compute visual tasks [1–5]. Many of these tasks require fast processing of real-time inputs. In autonomous navigation, for example, a network must be capable of identify obstacles with different textures and lighting conditions in real time. While CNNs are instrumental in providing high accuracy for such tasks, the time that it takes for the input to propagate through the trained network (forward propagation time), is large and precludes real-time operations. The main reason for such inefficiency is the computational complexity of CNNs, which is  $\mathcal{O}(HWk^2)$ , where  $H(W)$  is the height (width) of an image frame and  $k \times k$  is the size of the convolutional kernel. The challenge of enhancing the computational performance has driven significant development of hardware that is dedicated to computing convolutions, with graphics processing units (GPUs) and tensor processing units (TPUs) that deliver an order of magnitude acceleration. Even with such dedicated hardware, however, effective computation times are still sub-optimal and become too large for many applications, especially when high-resolution inputs are processed. For example, when applied to ResNet-18 for autonomous navigation, the Jetson Nano, NVIDIA hardware development kit, achieves at most 5 frames per second (fps) rate for images of dimensions  $1000 \times 1000$  [6]. This thesis investigates ways to increase the forward propagation speed of CNNs through computational and optical approaches. It aims to increase the performance of forward propagation runtime by updating parameters in spectral domain and further accelerate those networks by implementing the  $4f$  optical structure.

Since convolution is the most time-consuming operation, a possible gain in computational efficiency can be achieved by implementing the CNN in the Fourier domain (spectral domain) [7, 8]. The Fast Fourier Transform (FFT) operation has the complexity of  $\mathcal{O}(HW \log(HW))$  for the images

and the kernels. In the Fourier domain, the convolution is transformed into an elementwise product with only  $\mathcal{O}(HW)$  operations. While promising, in practice, the approach does not boost the forward propagation time. Due to nonlinearities that follow most of the layers in a CNN, the spectral approach ends up including a large number of costly forward and inverse Fourier transforms between successive layers. Optimizations of network architectures to be compatible with spectral operations were proposed, such as FCNN and Clebsch–Gordan Nets [9, 10]. These architectures, however, appear to suffer from a reduction in classification accuracy for complex visual tasks, for example, FCNN reaches less than half of state-of-the-art accuracy on CIFAR-10. Another branch of research developed in parallel is the introduction of spectral pooling layers for the purpose of adapting the spatially applied Max-Pooling operation to the spectral domain [11–13]. Unfortunately, these networks still entail nonlinear activations that require conversions between the spatial and spectral domains. Removing these nonlinear functions dramatically reduces the computational complexity of spectral CNNs but this often comes at the cost of reduced performance.

Another promising approach for accelerating computation is the development of optical neural networks (ONNs) that could replace electronic hardware [14]. ONNs utilize the inherent parallelism of light, which enables passive manipulation of massive amounts of data at the speed of light [15–18]. ONNs offer computation times that are nearly instantaneous in comparison to those provided by the best electronic hardware available. Specifically, a lens can perform a Fourier transform with  $\mathcal{O}(1)$  complexity [19]. In recent years, this promise has resulted in various ONNs for MNIST classification based on a sequence of diffractive masks in the Terahertz regime [17], vector-matrix multiplication using integrated photonics [20, 21], and hybrid optical-electronic networks leveraging the inherent Fourier transform property of lenses exploited in a  $4f$  architecture [18]. While impressive in their own right, the demonstrated ONNs to date have been limited in terms of the complexity of the scenes on which they operate, which have mostly been limited to low-resolution images with simple features like MNIST digits. When applied to more complex scenes (e.g., HIGH-10), these implementations exhibit low classification accuracy.

A major contributing factor to these limitations is that these networks lack the optical implementation of a nonlinear activation and were mostly constrained to linear operations in the optical

domain [18]. Achieving an optical nonlinearity requires very large optical power, a high-quality factor resonator, exotic materials, or a combination thereof. While some nonlinear functions are readily available, such as the square operation imparted by a detector, it is unclear how effective it is for achieving comparable performance to that of the ReLU nonlinearity, which is the most common and effective nonlinear activation for CNNs [22, 23]. If the nonlinearity is implemented electronically as a subsequent layer, the electronic signal needs to be converted back into the optical domain to enable additional optical processing. Such repeated signal transductions significantly increase the power consumption and latency, thus obviating the benefits of an ONN versus a more traditional electronic implementation [18].

To compensate for the lack of nonlinearity in an ONN, as well as to find an optimal level of performance for a fully linear spectral network, we develop a Knowledge Distillation (KD) training methodology to transfer the information from a nonlinear network (teacher) to a linear counterpart, SCLC (student). Originally, KD training was introduced for pruning of networks, i.e., knowledge from a large teacher network is transferred to a less complex student model [24, 25]. Trained with the KD approach, the student network typically converges faster and obtains better performance than it would achieve without the KD training. A common example for successful KD training is object classification in images. In this problem, the teacher classifies images and provides "soft labels" to the student during training along with the actual labels. A Kullback-Leibler (KL) divergence loss between the soft labels and the student model predictions is then optimized to take into account the teacher's predictions [26]. KD training is a generic approach and was applied to a variety of problems such as semantic segmentation in which soft labels are used for classification of each pixel [27, 28].

In this thesis, we adapted the KD training framework to circumvent the need for a nonlinearity by "distilling the nonlinearity" from the nonlinear, teacher CNN, to the linear counterpart SCLC. We find that for boosting student accuracy, the teacher and the student networks are required to be as architecturally similar as possible. The KD approach allows the student network to achieve state-of-art performance, exceeding that of previous training methods or networks in the spectral domain, and is also easily amenable to optical implementation because light propa-

gation is naturally described in Fourier space. The adapted KD training enables us to design a hybrid optical-electronic architecture for the SCLC network, where the optics serve as a linear frontend processing unit connected to an electronic backend that typically includes the last layer that corresponds to the task. We train this student network with KD training and demonstrate the performance of such a network on two common problems in which CNNs are leading computational methods: object classification and object segmentation . We show that the KD-trained SCLC can achieve performance easily surpassing that of a linear network trained with a standard training approach and nearing the performance of the nonlinear network.

## Chapter 2

### METHODOLOGY

#### 2.1 Knowledge Distillation (KD)

Knowledge Distillation is a machine learning method introduced for compression of neural networks [24]. In particular, it defines the transfer of knowledge from a large neural network model, called the *teacher*, to a small model, called the *student*. We show the schematic flow of KD in Fig. 2.1. KD assumes that the teacher model is already trained and performs the given task with high accuracy. Then the student model is trained with both the conventional approach of minimization of the loss between the model prediction and the training data (*Student loss*) and, in addition, compares the outcome of the student model with the teacher model through a temperature loss (*Temp loss*) using the KL divergence as the loss [4]. The optimization of both losses is performed through the back-propagation, which adjusts the parameters of the student model using Stochastic Gradient Descent [29], or ADAM [30] optimization.

It was shown that student models trained with a KD approach can achieve significantly higher accuracy than the same model trained on the data alone without a teacher network [26]. The advantage of KD stems from the proposal to compute the temperature loss which treats the prediction of the teacher model as ‘soft labels’ that inform the student model. In particular, conventional learning considers exclusively the labels in the training data as ‘hard labels’ and for tasks such as classification computes the probabilities distribution vector  $p^{hl}$ , where each element  $p_i^{hl}$  in the vector corresponds to the probability of the current input belonging to the class  $i$ . The softmax function is used to compute probabilities

$$p_i^{hl} = \exp(z_i) / \sum_j \exp(z_j). \quad (2.1)$$

where  $z_i$  is the student logits after the last fully connected layer. Training with hard labels is

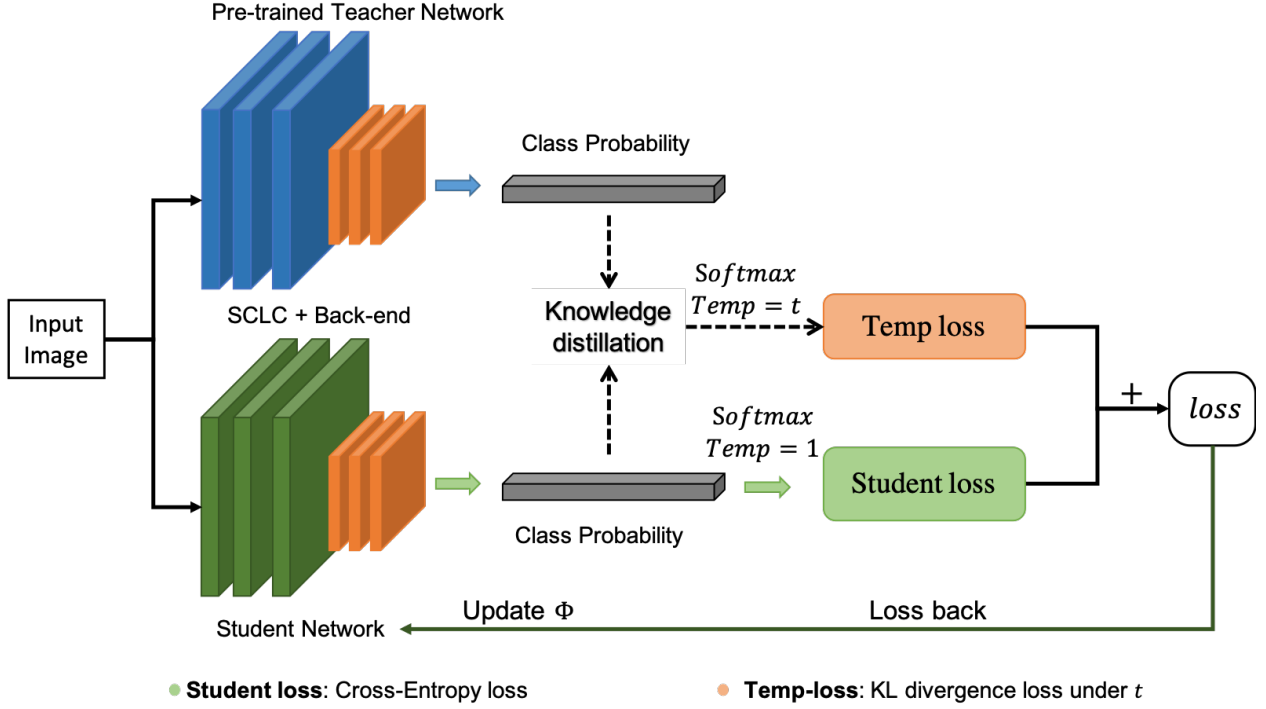


Figure 2.1: Illustration of KD training. Student network (green-bottom) parameters  $\Phi$  are updated according to an interpolation of two losses: the Student loss (Crossentropy loss between data and the student model output) and Temp loss (Cross-entropy loss or KL divergence) between the teacher network (blue-top) and student class distribution with a temperature parameter  $T$ .

a sensitive process, especially for compact networks such as the student model. This typically results in inefficient networks and convergence to poor local optima. To overcome this limitation, KD proposes to add soft labels generated by the teacher model. These labels are the probabilities that the teacher model generates. In particular, for each input, at the same time of computing  $p^{hl}$  with the student model, KD computes the soft probabilities vector,  $p^{sl}$ , with the teacher model according to

$$p_i^{sl} = \exp(y_i/T) / \sum_j \exp(y_j/T). \quad (2.2)$$

where,  $y_i$  are the logits of the teacher after the last fully connected layer. Through  $p^{sl}$  the teacher model contributes the probability estimates to the student model, knowledge that is unavailable from the data alone and is helpful for the student as providing extra information of the similarities

between classes. The similarities are important since these indicate the effective knowledge of the teacher model and subsequently assists in achieving a similar knowledge and performance in the student model. The two probability distributions  $p^{sl}$  and  $p^{hl}$  are taken into account (as an interpolation) to compute the overall loss used in training of the student model.

The probabilities contributed by the teacher are defined as soft since the softmax function has a softening parameter,  $T$ , named as the distillation temperature. The success of KD depends on the choice of  $T$ . In a well-trained teacher model, the correct class has a much higher probability than other classes. For a low value of  $T$ , the probability of the correct class will approach 1 while probabilities of other classes will be negligible and will not influence the training, due to  $p^{sl}$  being similar to  $p^{hl}$ . On the other hand, when  $T$  is too high,  $p_i^{sl}$  will approach a uniformly distributed vector and finally loses the distinction between the correct and incorrect classes. It is therefore important to choose  $T$  in between these two extreme cases such that  $p^{sl}$  would pass the similarities detected by the teacher to the student.

In practice, the distillation of knowledge from the teacher to the student is particularly effective when the differences in the overall architecture between the two networks are minimal. In network compression, it is typically the case that the architectural blocks are kept the same and the only change that is implemented is the reduction of the number of neurons in each block. This leads us to the proposition of the implementation of KD between nonlinear and linear CNNs, where besides exclusion of nonlinear components, the linear network will be kept as similar as possible to the nonlinear one.

## **2.2 The Spectral CNN Linear Counterpart (SCLC)**

We propose to construct a Spectral CNN Linear Counterpart (SCLC), the “Student Network”, to a spectral nonlinear CNN, the “Teacher Network”, based on an optical  $4f$  architecture. The “Teacher Network” takes the shape of a common CNN designed for generic tasks, for example, image classification or object segmentation. In the case of the classification, we consider CNNs that are structured with multiple layers of repeating operations of Convolution, Nonlinearity (ReLU), and Max-Pool (selecting the maximal value), as demonstrated in Fig. 2.2. In the case of object

segmentation, we consider a CNN of a “U” shape [31], where the input passes through similar operations of Convolution, ReLU, and Max-Pool, and in addition, the output of each layer contracts the input dimension up to a “bottle-neck” layer from which the representation is expanded with a set of inverse operations such as Up-Convolution, ReLU, and Max-Pool. To measure the performance of the proposed architectures, we concatenate them with a backend that corresponds to either a classification task (softmax fully connected backend layers) or a segmentation task (sequence of Up-Convolution backend layers).

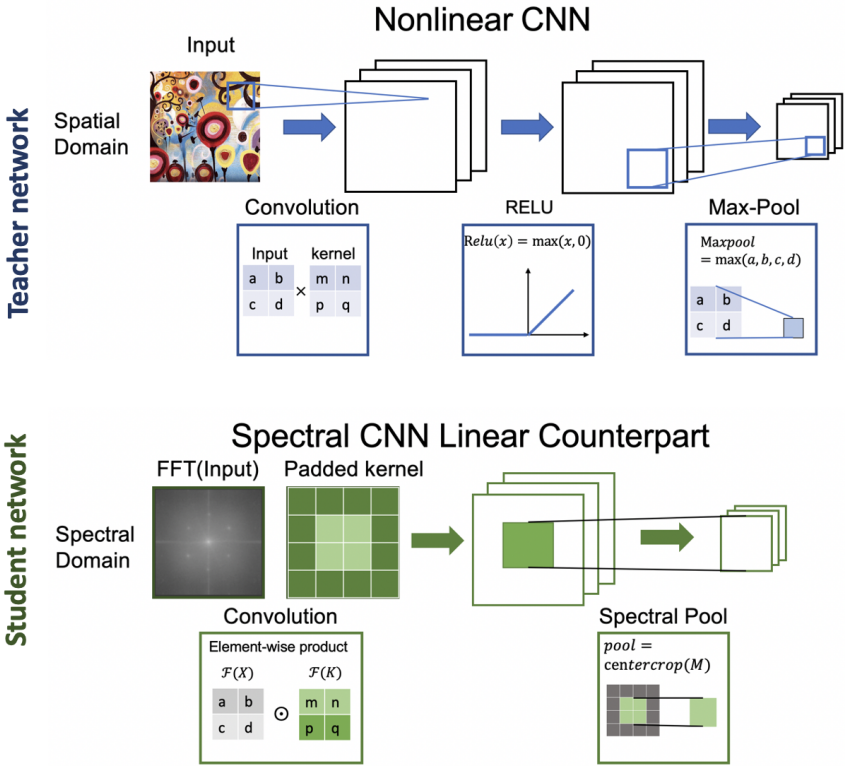


Figure 2.2: Nonlinear CNN (top) and the proposed substitute, Spectral CNN Linear Counterpart (SCLC) (bottom). Top: An exemplar nonlinear CNN with layers that include operations of Convolution, Nonlinear RELU activation, and Max-Pool. Bottom: SCLC of the CNN shown in top row. The Convolution corresponds to the elementwise product in the spectral domain. The RELU operation is excluded. The Max-Pool layer is represented by a center-crop operation in the spectral domain.

To obtain a network model that is realizable with optical components, the network requires a

conversion to the spectral domain, such that the input into the model is the Fourier transform of the input image and the operations such as convolution and pooling are implemented in the spectral domain [18]. Furthermore, the model should not include nonlinearities since computing those will require an inverse transform and transition from optical to electronic components. We describe the building blocks of such a proposed model below.

### 2.2.1 Convolution in SCLC

The convolution is performed as the elementwise product between the input images represented in the spectral domain and the kernels padded to the same dimensions as the inputs. The input into the convolution is a four-dimensional tensor (S: batch size, C: number of input channel, H: height, W: width), and the 2D convolution operation with a stride of 1 is calculated via

$$y(i, j) = x * k = \sum_{h=0}^{H-1} \sum_{w=0}^{W-1} x(h, w) \times k(i-h, j-w) \quad (2.3)$$

where  $x$  is the input image,  $k$  is the kernel and  $*$  indicates the Convolution operation. We implement the convolution in SCLC in the spectral domain by using an FFT and an elementwise product

$$Y = \mathcal{F}(x * k) = \mathcal{F}(x) \odot \mathcal{F}(k) \quad (2.4)$$

The function  $\mathcal{F}$  denotes the FFT operation, and  $\odot$  indicates the elementwise product, which requires the input and the kernels to have the same dimension  $H \times W$ . During training, the spectral convolution kernel is updated according to

$$\begin{aligned} \sigma_X &= \nabla_X \mathcal{L}|_{X=X_0} = \sigma_Y \odot K_0 \\ \delta_k &= \nabla_K \mathcal{L}|_{K=\mathcal{F}(k_0)} = \sigma_Y \odot X_0 \\ k_1 &= k_0 + \lambda [\mathcal{F}^{-1}(\delta_K)], \end{aligned} \quad (2.5)$$

where  $\sigma_X(\sigma_Y)$  is the error from the previous (next) layer;  $X_0$  and  $K_0$  are, respectively, the input and kernel in the forward propagation;  $\nabla_K(\nabla_X)$  is the gradient operator w.r.t. the kernel (input); and  $\mathcal{L}$  is the loss function. The updates to the elements in the kernel are computed by applying an inverse FFT (i-FFT) and multiplying by the learning rate  $\lambda$ .

Once the network has been trained, the implementation with spectral convolution components has significant benefits for inference latency for a given input (forward propagation). Indeed, assuming an input image size of  $(H, W)$  and square kernel size of  $(k, k)$ , the complexity of one spatial convolution in the image domain is  $\mathcal{O}(HWk^2)$ . On the other hand, the spectral convolution composed of elementwise products is of complexity  $\mathcal{O}(HW)$ . In the optical setup, elementwise operations can be performed in parallel and thus the runtime, being independent of  $H$  and  $W$ , is further reduced to  $\mathcal{O}(1)$ . In addition, the computational complexity of an FFT is  $\mathcal{O}(HW \log(HW))$ , which brings the overall complexity for the spectral convolution to be  $\mathcal{O}(HW \log(HW))$  [10,32,33]. As the proposed network is in the spectral domain, the FFT and i-FFT transforms are needed to be applied at the beginning and at the end of the network for images only, such that the complexity of the intermediate layers will only apply FFT to kernels, which greatly reduces the transformation times. Additionally, the main acceleration occurs in the optical implementation since the spectral transforms can be achieved through phase transforming components which operate at the speed of light. The combination of convolution and transform in optical domain would correspond to instantaneous running time of  $\mathcal{O}(1)$ , when compare to electronic running time for all layers of the network.

### 2.2.2 Pooling in SCLC

Implementation of pooling in the spectral domain needs to take into consideration both mimicking the effect of the standard Max-Pool used in CNNs and to be a practical implementation with optical components. In a nonlinear CNN, common pooling operations are Average-Pooling or Max-Pooling, which select the average of the elements or the maximum element, respectively, from the region of the feature map covered by the pooling filter shown in Fig. 2.2. The purpose of the pooling layers is to reduce the dimensions of the feature maps and to find the representative elements to be transferred forward from the feature map. We propose to substitute the Max-Pool with a different pooling function, called ‘‘Spectral pool’’, which will enable similar functionality in the spectral domain. The Fig. 2.3 demonstrated the advantage and effect of the spectral pooling for different output dimensions. In particular, we propose to replace the pooling layer with a lin-

Table 2.1: Comparison of forward propagation computational complexity between Nonlinear CNN, SCLC implemented on electronic hardware, SCLC implemented on optical hardware

Structure	Nonlinear CNN	SCLC (Computational)	OSCLC (Optical)
Convolution layers	$\mathcal{O}(HWk^2)$	Spatial: $\mathcal{O}(HW \log(HW))$ Spectral: $\mathcal{O}(HW \log(HW))$	$\mathcal{O}(1)$
Pooling layers	$\mathcal{O}(HWk^2)$	Spatial: $\mathcal{O}(k^2 \log(k^2))$ Spectral: $\mathcal{O}(1)$	$\mathcal{O}(1)$

ear, low-pass filter in the spectral domain. The forward and backward propagation in the layer are defined as

$$y = \mathcal{F}^{-1}(\text{CROP}(\mathcal{F}(x), (H', W'))) \quad (2.6)$$

$$z = \mathcal{F}^{-1}(\text{PAD}(\mathcal{F}(y/x^*), (H, W))) \quad (2.7)$$

where the *CROP* operation in forward propagating keeps the center part of the spectral feature map and reduce the total size from  $(H, W)$  to  $(H', W')$ . The *PAD* operation in backward propagation matches the dimension of gradients outputs from  $(H', W')$  to  $(H, W)$  by padding zeros instead of cropped elements. The  $\mathcal{F}$  and  $\mathcal{F}^{-1}$  are applied if the networks are in spatial domain. Those operations are similar to max pooling since the idea of Max-Pool is to find the representative element in each kernel, while it goes over the feature map and effectively selects the most descriptive features. Spectral-Pool drops the high frequencies in the Fourier domain such that it achieves a less noisy output and captures the dominant features, similar to the effect of the conventional Max-Pool operation. It was shown that training with spectral pooling has better convergence properties compared to Max-Pooling [11]. Furthermore, a spectral pooling implementation is of lower computational complexity and is furthermore inherently amenable to optical implementation as light is well described in the Fourier domain  $\mathcal{O}(1)$  complexity, see Table 2.1.

Typically, an activation function, such as ReLU, tanh, or  $\sigma$  is applied to extract the features after convolution to avoid extremities in neural units values and maintain network stability. These nonlinearities, however, are nearly impossible to achieve in a practical optical implementation.




Figure 2.3: Pooling in SCLC: Comparison for different pooling schemes in different dimension reductions. Top: Max pooling; Bottom: Spectral Pooling. The Spectral pooling retains more information on any arbitrary output dimension.

Thus, the SCLC skips the nonlinear activation function to make the network structure amenable to optical implementation. The lack of nonlinearity will be circumvented by the application of KD to maximize the performance achieved using linear operations only.

### 2.3 Knowledge Distillation in SCLC

KD training requires a teacher and student model both performing the same task and exhibiting similarity in their architectures. In previous applications, KD training was applied to model pruning, where the teacher and the student models have exactly the same structure with the student having a fraction of the units of the teacher. Here, we extend KD application and consider the teacher and the student models with a similar number of units; however, the student is an adapted version of the teacher operations lacking the nonlinear activation and has a revised pooling operator. In particular, the teacher model is a pre-trained nonlinear CNN, which is a state-of-the-art system for that particular task. The student network is the SCLC that corresponds to that system with the architectural changes described in the previous sections. The objective of KD training is to optimize the total loss subject to updating the weights,  $\Phi$ , of the SCLC student model only. The

total loss is a linear combination of the Temperature loss and the Student loss. The weights are updated according to implementation of back-propagation that optimizes the total loss.

We select the Temperature loss to be the KL function between the soft labels ( $p^{sl,t}$ ) from the pre-trained nonlinear CNN model (*teacher*) and predictions ( $p^{sl,s}$ ) from SCLC (*student*) both distilled with the same temperature  $T$ . We chose KL loss over cross entropy because it includes an extra penalty on the direction of the loss, which facilitates convergence. The student loss is the standard cross-entropy loss between the data labels and SCLC probabilities ( $p^{hl}$ ). The total loss is then calculated as a weighted summation of the two losses

$$Loss(x, \Phi) = \alpha \mathcal{L}_C(y, p^{hl}) + (1 - \alpha) \mathcal{L}_K((p^{sl,t}; T = \tau), (p^{sl,s}; T = \tau)), \quad (2.8)$$

where  $x$  corresponds to the input,  $y$  is the training data,  $\Phi$  are student model weights,  $\mathcal{L}_C$  is the cross-entropy loss function,  $\mathcal{L}_K$  is the KL divergence loss function,  $p^{hl}$  corresponds to the student model hard predictions,  $p^{sl,s}$  corresponds to the student predictions under given teacher model probabilities  $p^{sl,t}$ , and  $\alpha$  is the weighting parameter.

#### 2.4 Optical Implementation of SCLC

In this section, we explore how the proposed SCLC network can be implemented using free-space optics, i.e. Optical SCLC (OSCLC). The choice of free-space optics is motivated by the large number of information channels available to us in such an implementation [23]. There are three components of the OSCLC that need to be considered for optical implementation: convolutional layers, spectral pooling, and summation of different channels. The Fig.2.4 demonstrates the schematic of a single optical  $4f$  correlator structure.

Each convolutional layer can be implemented by a  $4f$  correlator architecture to further accelerate the forward propagation speed [15, 34–36]. A typical  $4f$  correlator architecture comprises two lenses of equal focal length spaced apart at  $2f$  distance and with input and output planes located in the front and the back focal planes of the first and second lenses respectively. The first lens produces a Fourier transform of the input scene at the focal plane. A mask placed in that focal plane provides the point-wise multiplication implementing the convolution and the second lens performs




Figure 2.4: Optical  $4f$  structure. The OSCLC can be implemented in  $4f$  structure parallelly.

the inverse Fourier transform. Therefore, a  $4f$  correlator architecture is able to work as an equivalent architecture for a single channel of a linear spectral CNN counterpart. The computational complexity can decrease to  $\mathcal{O}(1)$  for any operations in spectral domain and will not be exponential growth with the input image resolution/ pixels. The spectral pooling in the OSCLC can also be implemented by a low-pass filter, where we block high frequency components in the Fourier plane.

The summation of different channels is a basic principle in deep neural networks and is widely used in object classification and segmentation. If there are no summations after elementwise products, the channels will grow exponentially, requiring a massive number of kernels after only a few convolutional layers, which makes it impossible for real optical implementations. Although fewer kernels in the convolution layers can alleviate this challenge, the overall accuracy, especially in complex scenarios, would suffer. Such a summation can be implemented using various techniques used for coherent beam combining [37, 38]. We note that, while free-space optics tends to be bulky and prone to misalignment, recent demonstrations of meta-optics and volume optics [39, 40] exhibit complicated free-space optics in a compact form factor, possibly in a monolithic fashion mitigating any misalignment.

## Chapter 3

### COMPUTATIONAL EXPERIMENTS

We evaluate our proposed SCLC on two distinct tasks involving image inputs: (i) object classification and (ii) object segmentation. Object classification task associates each image with a class which corresponds to the object photographed in the image, for example, the task could be to classify whether the image depicts a cat, a dog, or a car. Object segmentation is a pixel-level classification task. It aims to determine for each pixel in a given image to which class it belongs. Effectively, object segmentation partitions the input image into regions that correspond to classes. For example, in the case of a photograph of a person, regions would correspond to the background, face, hair, etc. Our evaluation consists of determining the accuracy and the forward propagation runtime that the SCLC and OSCLC achieve compared to a standard CNN.

We test the performance metrics on a variety of common benchmarks for each task. For object classification, we perform experiments on (i) Kaggle Cats and Dogs Challenge [41], (ii) Cifar-10 [42], and (iii) HIGH-10 (a subset of ImageNet) [43]. For object segmentation, we perform experiments on (i) Kaggle’s Carvana Image Masking Challenge (Cars Segmentation) [44], (ii) Face Recognition [45] and (iii) VOC2012 [46]. Training of all networks is performed on a Tesla P100 GPU with the Google Colab platform. For all networks, the initial learning rate for student networks is 0.0001 with momentum 0.9 and weight decay 0.0005. We set the batch size according to the available memory: 16 for the classification task and 1 for the segmentation task. We use the cross-entropy loss to calculate the student loss, and KL divergence to calculate the Temp loss.

#### ***3.1 Spectral Counterpart Layer Efficiency***

We first investigate the runtime of forward propagation when variable input dimensions ( $32 \times 32$  to  $1024 \times 1024$ ) are considered for AlexNet (nonlinear CNN) vs. SCLC, and show our results

in Fig. 3.1. We evaluate forward propagation time for convolutional layers and pooling layers separately. Forward propagation time of convolutional layers in AlexNet grows exponentially with input size. The runtime of SCLC grows exponentially as well but with a significantly slower rate. For example, For High Definition (HD) input size image, the runtime of SCLC convolution corresponds to 100ms vs. 500ms for AlexNet convolution. For pooling layers, AlexNet forward propagation runtime grows exponentially with input size with a smaller rate than the convolution. For SCLC, due to the pooling being spectral, the runtime is constant and input independent.




Figure 3.1: Comparison of forward propagation time between AlexNet (blue line) and its SCLC (red line). Left: Propagation time through a single convolutional layer for variable image dimensions; Right: Propagation time through a single pooling layer for variable image dimensions. The dashed lines mark the resolution of an HD image and its corresponding runtime in AlexNet and SCLC.

### 3.2 Classification

We evaluate our proposed SCLC on different classification datasets to estimate the accuracy that SCLC can achieve when trained with KD. We compare the accuracy to that of the teacher network, AlexNet. The first dataset that we consider is the Kaggle Cats and Dogs Challenge [41] which consists of 125000 images of dimensions  $96 \times 96$  associated with two classes: a cat or a dog. Additional dataset that we consider is Cifar-10 [42] which consists of 60000 images of dimension of  $32 \times 32$  including 10 classes. The third dataset that we consider is High-10 which is a subset

of ImageNet [43] and consists of approximately 10000 annotated images, equally distributed in 10 classes with images of resolution  $500 \times 300$ . During training of AlexNet all images are resized or cropped to dimensions of  $224 \times 224$  to match AlexNet’s input size. To evaluate the performance of SCLC in classification we add a single fully connected backend layer to it. The backend layer is intended to be implemented with electronics in OSCLC since includes nonlinearity. We employ the KD approach to train SCLC with AlexNet being the teacher network and when the training converges we test SCLC variants against AlexNet.

In particular, we compare the SCLC trained with and without the KD approach or a variation of AlexNet with square nonlinearity that could be realized optically with 4f components. We show the results of the comparison for the three benchmarks in Table 3.1. We observe that for all benchmarks, KD training significantly enhances the accuracy of the classification achieved by the SCLC (by 12.5% on average). Indeed, KD contribution appears to be essential in generating an SCLC network with robust accuracy. On Kaggle’s Cats and Dogs classification SCLC achieves 90.60% (6% below the accuracy of AlexNet), on Cifar-10 classification it achieves 80.80% (5% below the accuracy of AlexNet) and on HIGH-10 classification it achieves 81.4% ( $\approx$  12% below the accuracy of AlexNet). These results are encouraging since SCLC trained with standard training has a much bigger gap of 16%, 20%, 23% between its accuracy and the accuracy of AlexNet. The KD approach appears to close this gap by more than half for all benchmarks. Furthermore, when the RELU nonlinearity is modified to square nonlinearity (SQ-AlexNet) both KD and standard training do not result with sufficiently accurate network. This is because the square nonlinearity will magnify the parameters of the model. This observation indicates that the KD approach is effective in regimes where the network architecture of the student is kept as close to the teacher as possible.

Notably, to match AlexNet input dimensions, all experiments were implemented with the same input image resolution of  $224 \times 224$ . We therefore explore with HIGH-10 dataset (that includes higher resolution images) how accuracy varies if the resolution of the input is increased. For each input image resolution down sampling to 224, the teacher (AlexNet) is first trained and then the student (SCLC/OSCLC) is trained under the supervision of the teacher model. We show in Fig. 3.2

that increasing the resolution of the input increases the accuracy of both the teacher and student models (with a linear rate). Consequently, SCLC trained with higher resolution inputs can surpass the accuracy of the teacher network trained with lower resolution dataset, e.g., SCLC with input of  $2^7 \times 2^7$  dimensions performs similarly ( $\approx 80\%$ ) to the teacher with input of  $2^6 \times 2^6$  dimensions. While the accuracy of the two is similar, the runtime of SCLC is expected to be more favorable since teacher’s runtime grows exponentially. Indeed, we show that for 80% accuracy both SCLC and OSCLC are at least 5 times faster than ( $\approx 5ms$  and  $\approx 1ms$ ) than AlexNet ( $\approx 25ms$ ).

We further demonstrate the efficient runtime of SCLC compared to AlexNet in Table 3.1 columns 3,4 for input size of  $224 \times 224$ . While AlexNet runtime for a single image is  $350ms$ , SCLC runtime drops by an order of magnitude to  $11ms$  for a single image. Simulated runtime of OSCLC drops by another order of magnitude to  $0.6ms$ . The estimation of OSCLC forward propagation runtime includes three parts: 1) running time of the optical structure, 2) transduction time between optics and electronics, and 3) running time of the electronic backend. Since light propagation is very short,  $\approx ps$ , the main contribution comes from the signal transduction ( $\approx 0.32ms$  for 100kb images via USB 3.0 protocol at a rate of 2500Mbit/s) and the backend propagation time ( $\approx 0.28ms$  for a single image on a GPU (Tesla P100)).




Figure 3.2: Accuracy and Forward propagation time for classification task on High-10 dataset. Left: AlexNet and SCLC accuracy for varying input resolution; Right: Forward propagation runtime of AlexNet, SCLC, and OSCLC for varying input resolution.

Table 3.1: Accuracy and estimate the efficiency in forward propagation in classification

Dataset	Model(input:224 × 224)	Accuracy	Runtime (ms/img)	Optical Runtime (ms/img)
Cats vs. Dogs	AlexNet	96.10%	350.66	-
	SCLC	79.50%	11.05	0.61
	SCLC + KD	90.60% (+11.10%)	11.05	0.61
	SQ-AlexNet	51.70%	12.04	0.61
	SQ-AlexNet+ KD	51.72%	12.04	0.61
Cifar-10	AlexNet	85.09%	350.66	-
	SCLC	65.45%	11.05	0.61
	SCLC+KD	80.80% (+15.35%)	11.05	0.61
High-10	AlexNet	93.95%	350.66	-
	SCLC	70.12%	11.05	0.61
	SCLC + KD	81.40% (+11.28%)	11.05	0.61

To demonstrate how different components contribute to SCLC performance, we investigated ablations of the SCLC architecture (Table 3.2). We first compared Spectral Pooling with Max Pooling in the SCLC. We find that Spectral Pooling corresponds to faster convergence and to 1.71% higher accuracy than Max Pooling. In addition, we test the contribution of the frontend (SCLC), the backend and the KD training of them. When only the backend is considered, the accuracy drops to 41%. Adding the frontend to it and training the model on the same dataset corresponds to an increase of  $\approx 30\%$  such that the accuracy becomes 70.12%. Further application of KD training of both the frontend and the backend boosts the accuracy and convergence, such that the accuracy is enhanced by another  $\approx 10\%$  to 81.4%.

### 3.3 Object Segmentation

In addition to image classification, we explore the construction of an SCLC for object segmentation as shown in Fig. 3.3. For such a task, a standard CNN is of a U shape with a sequence of convolutions and up convolutions. We choose the teacher network to be a U-Net [31]. KD training

Table 3.2: Ablation studies of pooling components, SCLC, backend and KD training

Structure	Accuracy
Max Pooling	68.41%
Spectral pooling	70.12%(+1.71%)
Backend only	41.40%
SCLC (Frontend) + Backend	70.12% (+28.72%)
SCLC (Frontend) + Backend + KD	81.40% (+11.28%)

corresponds to pixel-level loss for both soft predictions and soft labels. The frontend for this task is the sequence of contracting convolutions while the backend corresponds to a sequence of up convolutions.

The first benchmark that we consider is the Kaggle’s Carvana Image Masking Challenge that consists of 5088 cars with an original resolution of  $1920 \times 1280$ . The dataset is randomly split into 4580 and 508 images for training and testing, respectively. Our second benchmark is the Face Recognition dataset, which consists of 2000 images with 1700 for training and 300 for testing. Our third benchmark is VOC2012, which consists of 2913 images of original resolution  $500 \times 375$ , which includes 20 classes and one background class. The images in all datasets are down-sampled to  $960 \times 640$  or smaller due to limitations of GPU memory.

Our results are shown in Table 3.3. For both Car and Face datasets we observe that the SCLC performs relatively well and obtains accuracy that falls from that of the teacher U-Net by only a few percent. We observe that consideration of higher resolution inputs corresponds to enhanced accuracy in classification experiments. For the VOC2012 dataset, which includes lower resolution images and has more segmentation classes, both the teacher and SCLC perform with rather low accuracy below 80%. This example demonstrates that in such problems, it is crucial to consider the fully available image resolution. We compare the computational efficiency of the SCLC against the teacher in terms of frames per second rate and find that the SCLC is approximately  $2 \times$  faster than U-Net. While such a speedup is more modest than in the image classification task, the rate

Table 3.3: Accuracy and forward propagation rate (frames per second) for object segmentation task tested on three benchmarks.

Dataset	Network	Accuracy	Rate (higher is better)
Car Segmentation	U-Net	98.02%	7.36 fps
	OSCLC	97.1%	12.7 fps
Face Recognition	U-Net	95.58%	9.38fps
	OSCLC	91.39%	15.6fps
VOC2012	U-Net	75.51%	7.36fps
	OSCLC	62.21%	12.7fps

is closer to a real-time operation rate (30 fps) or alternatively allows for consideration of larger input images that may correspond to enhanced segmentation with the same frame rate as that of the teacher. The main contribution to SCLC runtime is the electronic backend that consists of up-convolution layers, which have a computational complexity of  $\mathcal{O}(HWk^2)$  (compared to image classification backend, which is a fully connected single layer).

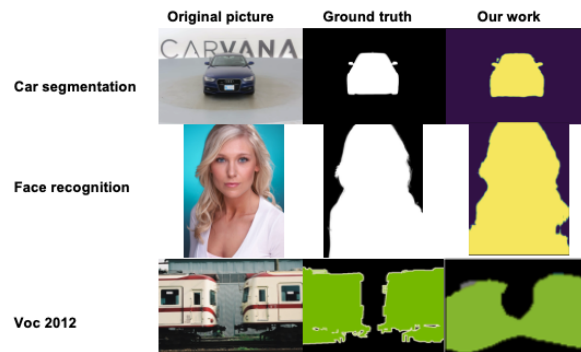


Figure 3.3: Examples of object segmentation for the considered benchmarks of (i) Car Segmentation, (ii) Face Recognition, and (iii) VOC2012. Left to right: Input image from datasets, Ground truth, Our work (SCLC/ OSCLC).

## Chapter 4

### CONCLUSION

Over the last decade, various CNNs have been developed and deployed as robust systems for ubiquitous visual tasks. While such incorporation in multiple applications is remarkable, scaling up CNNs for tasks that require high speed and energy efficiency entail significant challenges. This is due to the fact that the computational runtime of convolutional layers increases exponentially with input size. Such a constraint eventually leads to CNNs exceeding the limit of power consumption and as such they limit the accuracy and rate at which the system can perform. It is often the case that these limitations prohibit real-time deployment of CNNs. Representation of CNNs in the spectral domain in conjunction with an optical platform that supports parallel, elementwise product computations has the potential to overcome the exponential increase in computation runtime. However, since transitions from optical to electronic systems are computationally ineffective, it is required to compute the additional, typically nonlinear, layers of CNNs with the same optical platform. Performing such operations with optical components is currently impractical. To that end, we propose a spectral linear version of CNN (SCLC) that is also optically realizable (OSCLC) with 4f components and investigate whether it can serve as an effective counterpart for CNN. While it is indeed possible to design such a system, we show that the accuracy significantly drops when nonlinear activations are not present. To address the gap in accuracy between CNN and its SCLC, we propose a novel training approach based on Knowledge Distillation (KD). The approach assists in training the SCLC by involving the original CNN in the training as the teacher in addition to the standard optimization with the training data. We show that such training is promising in circumventing the need for a nonlinearity and represents a generic training procedure that could be applied to various CNNs and their respective SCLCs. In particular, we consider two visual tasks and six benchmarks and show that KD training applied to training a SCLC boosts its accuracy by

more than half of the gap created by exclusion of nonlinearities. We further show that increasing input dimensions has almost no impact on computational efficiency; however, increasing dimensions is able to further close the accuracy gap or even surpass the accuracy of CNN operating with low resolution inputs.

While our focus here is on realizing an optically viable architecture for an SCLC, it is notable that our analysis and experiments indicate a computational benefit for an SCLC trained with KD even for standard electronic components such as a GPU. For the OSCLC, we find that keeping the backend layer, which is implemented electronically, as minimal as possible is critical to fully make use of the benefits of optical parallelism and its speed. Indeed, as we demonstrate in our simulations, in the image classification task, in which the backend consists of a single layer, optical forward propagation reduces the runtime for standard input image dimensions by two orders of magnitude. In the object segmentation task, the backend is more complex. It consists of up-convolutions and thus while  $\times 2$  computational speedup is achieved, it is not as remarkable as the speedup in the case of the classification task. Furthermore, our investigation indicates that key components of CNNs, such as convolution and pooling, can indeed be adapted to the optical domain and for such adaptation the training process needs to be an integrated training procedure that involves both the original CNN, its adapted counterpart, and the electronic backend. We show that KD training presents a plausible integrated training procedure for these purposes. Adaptation of additional components alongside further development of KD-based training may pave the way toward optical-electronic deployment of CNNs to provide a further leap in enhancement of CNN performance for complex and real-time tasks.

## BIBLIOGRAPHY

- [1] E Kakkava, N Borhani, Y Pu, C Moser, and D Psaltis. Image classification and reconstruction through multimode fibers by deep neural networks. In *2018 Conference on Lasers and Electro-Optics Pacific Rim (CLEO-PR)*, pages 1–2. IEEE, 2018.
- [2] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [3] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [4] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [5] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [6] Sparsh Mittal. A survey on optimized implementation of deep learning models on the nvidia jetson platform. *Journal of Systems Architecture*, 97:428–442, 2019.
- [7] Michael Mathieu, Mikael Henaff, and Yann LeCun. Fast training of convolutional networks through ffts. *arXiv preprint arXiv:1312.5851*, 2013.
- [8] Nicolas Vasilache, Jeff Johnson, Michael Mathieu, Soumith Chintala, Serkan Piantino, and Yann LeCun. Fast convolutional nets with fbfft: A gpu performance evaluation. *arXiv preprint arXiv:1412.7580*, 2014.
- [9] Risi Kondor, Zhen Lin, and Shubhendu Trivedi. Clebsch–gordan nets: a fully fourier space spherical convolutional neural network. In *Advances in Neural Information Processing Systems*, pages 10117–10126, 2018.
- [10] Harry Pratt, Bryan Williams, Frans Coenen, and Yalin Zheng. Fcnn: Fourier convolutional neural networks. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 786–798. Springer, 2017.

- [11] Oren Rippel, Jasper Snoek, and Ryan P Adams. Spectral representations for convolutional neural networks. In *Advances in neural information processing systems*, pages 2449–2457, 2015.
- [12] Bochen Guan, Jinnian Zhang, William A Sethares, Richard Kijowski, and Fang Liu. Specnet: Spectral domain convolutional neural network. *arXiv preprint arXiv:1905.10915*, 2019.
- [13] Junxuan Li, Shaodi You, and Antonio Robles-Kelly. A frequency domain neural network for fast image super-resolution. In *2018 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2018.
- [14] Gordon Wetzstein, Aydogan Ozcan, Sylvain Gigan, Shanhui Fan, Dirk Englund, Marin Soljačić, Cornelia Denz, David AB Miller, and Demetri Psaltis. Inference in artificial intelligence with deep optics and photonics. *Nature*, 588(7836):39–47, 2020.
- [15] L Cutrona, EN Leith, C Palermo, and L Porcello. Optical data processing and filtering systems. *IRE Transactions on Information Theory*, 6(3):386–400, 1960.
- [16] Julian Bueno, Sheler Maktoobi, Luc Froehly, Ingo Fischer, Maxime Jacquot, Laurent Larger, and Daniel Brunner. Reinforcement learning in a large-scale photonic recurrent neural network. *Optica*, 5(6):756–760, 2018.
- [17] Xing Lin, Yair Rivenson, Nezh T Yardimci, Muhammed Veli, Yi Luo, Mona Jarrahi, and Aydogan Ozcan. All-optical machine learning using diffractive deep neural networks. *Science*, 361(6406):1004–1008, 2018.
- [18] Shane Colburn, Yi Chu, Eli Shilzerman, and Arka Majumdar. Optical frontend for a convolutional neural network. *Applied optics*, 58(12):3179–3186, 2019.
- [19] Joseph W Goodman. *Introduction to Fourier optics*. Roberts and Company Publishers, 2005.
- [20] J Feldmann, N Youngblood, M Karpov, H Gehring, X Li, M Stappers, M Le Gallo, X Fu, A Lukashchuk, AS Raja, et al. Parallel convolutional processing using an integrated photonic tensor core. *Nature*, 589(7840):52–58, 2021.
- [21] Yichen Shen, Nicholas C Harris, Scott Skirlo, Mihika Prabhu, Tom Baehr-Jones, Michael Hochberg, Xin Sun, Shijie Zhao, Hugo Larochelle, Dirk Englund, et al. Deep learning with coherent nanophotonic circuits. *Nature Photonics*, 11(7):441, 2017.
- [22] Dirk Englund, Arka Majumdar, Michal Bajcsy, Andrei Faraon, Pierre Petroff, and Jelena Vučković. Ultrafast photon-photon interaction in a strongly coupled quantum dot-cavity system. *Physical review letters*, 108(9):093604, 2012.

- [23] Albert Ryou, James Whitehead, Maksym Zhelyeznyakov, Paul Anderson, Cem Keskin, Michal Bajcsy, and Arka Majumdar. Free-space optical neural network based on thermal atomic nonlinearity. *arXiv preprint arXiv:2102.04464*, 2021.
- [24] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [25] Asit Mishra and Debbie Marr. Apprentice: Using knowledge distillation techniques to improve low-precision network accuracy. *arXiv preprint arXiv:1711.05852*, 2017.
- [26] Mary Phuong and Christoph Lampert. Towards understanding knowledge distillation. In *International Conference on Machine Learning*, pages 5142–5151, 2019.
- [27] Quanquan Li, Shengying Jin, and Junjie Yan. Mimicking very efficient network for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6356–6364, 2017.
- [28] Yifan Liu, Ke Chen, Chris Liu, Zengchang Qin, Zhenbo Luo, and Jingdong Wang. Structured knowledge distillation for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2604–2613, 2019.
- [29] Léon Bottou. Stochastic gradient descent tricks. In *Neural networks: Tricks of the trade*, pages 421–436. Springer, 2012.
- [30] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [31] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [32] Jae-Han Lee, Minhyeok Heo, Kyung-Rae Kim, and Chang-Su Kim. Single-image depth estimation based on fourier domain analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 330–339, 2018.
- [33] Armin Kappeler, Sushobhan Ghosh, Jason Holloway, Oliver Cossairt, and Aggelos Katsaggelos. Ptychnet: Cnn based fourier ptychography. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 1712–1716. IEEE, 2017.
- [34] Demetri Psaltis, David Brady, Xiang-Guang Gu, and Steven Lin. Holography in artificial neural networks. In *Landmark Papers On Photorefractive Nonlinear Optics*, pages 541–546. World Scientific, 1995.

- [35] Taiwei Lu, Shudong Wu, Xin Xu, and TS Francis. Two-dimensional programmable optical neural network. *Applied optics*, 28(22):4908–4913, 1989.
- [36] Demetri Psaltis, David Brady, and Kelvin Wagner. Adaptive optical networks using photorefractive crystals. *Applied Optics*, 27(9):1752–1759, 1988.
- [37] Tim D Gerke and Rafael Piestun. Aperiodic volume optics. *Nature photonics*, 4(3):188–193, 2010.
- [38] Alan Zhan, Shane Colburn, Christopher M Dodson, and Arka Majumdar. Metasurface freeform nanophotonics. *Scientific reports*, 7(1):1–9, 2017.
- [39] Maike Prossotowicz, Andreas Heimes, Daniel Flamm, Florian Jansen, Hans-Jürgen Otto, Aleksander Budnicki, Alexander Killi, and Uwe Morgner. Coherent beam combining with micro-lens arrays. *Optics Letters*, 45(24):6728–6731, 2020.
- [40] Ihsan Fsaifes, Louis Daniault, Severine Bellanger, Matthieu Veinhard, Jerome Bourderionnet, Christian Larat, Eric Lallier, Eric Durand, Arnaud Brignon, and Jean-Christophe Chanteloup. Coherent beam combining of 61 femtosecond fiber amplifiers. *Optics Express*, 28(14):20152–20161, 2020.
- [41] Kaggle dogs vs. cats redux: Kernels edition. <https://www.kaggle.com/c/dogs-vs-cats>, 2016.
- [42] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Technical Report*, 2009.
- [43] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012.
- [44] Carvana image masking challenge. <https://www.kaggle.com/c/carvana-image-masking-challenge>, 2017.
- [45] Xiaoyong Shen, Xin Tao, Hongyun Gao, Chao Zhou, and Jiaya Jia. Deep automatic portrait matting. In *European conference on computer vision*, pages 92–107. Springer, 2016.
- [46] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>, 2012.
- [47] Benedikt Pfülb and Alexander Gepperth. A comprehensive, application-oriented study of catastrophic forgetting in dnns. *arXiv preprint arXiv:1905.08101*, 2019.

- [48] Tyler Highlander and Andres Rodriguez. Very efficient training of convolutional neural networks using fast fourier transform and overlap-and-add. *arXiv preprint arXiv:1601.06815*, 2016.
- [49] James Tan, Zengguang Cheng, Xuan Li, Nathan Youngblood, Utku E Ali, C David Wright, Wolfram HP Pernice, and Harish Bhaskaran. Monadic pavlovian associative learning in a backpropagation-free photonic network. *arXiv preprint arXiv:2011.14709*, 2020.

## Appendix A

**APPENDIX RESULTS**

To give more comprehensive demonstration on this thesis, we offer the details of the network structure and more experiments results in the Appendix result.

**A.1 SCLC Model**

Table A.1: Implementation details of CNN, Optical fronted CNN and SCLC

Layer Number	CNN	Optical fronted CNN	SCLC
(0)	Conv2d(3, 64)	Optical Conv2d (3, 64)	Spec-Conv2d(3, 64)
(1)	ReLU	ReLU	-
(2)	MaxPool2d	Max-Pool2d	Spectral-Pool
(3)	Conv2d(64, 192)	Conv2d(64, 192)	Spec-Conv2d(64, 192)
(4)	ReLU	ReLU	-
(5)	MaxPool2d	MaxPool2d	Spectral-Pool
(6)	Conv2d(192, 256)	Conv2d(192, 256)	Spec-Conv2d(192, 256)
(7)	ReLU	ReLU	-
(8)	Conv2d(256, 256)	Conv2d(256, 256)	Conv2d(256, 256)
(9)	ReLU	ReLU	Optical-Electric convert
(10)	MaxPool2d	MaxPool2d	MaxPool2d
(Backend)	Fully connected backend	Fully connected backend	Fully connected backend
(Output)	LogSoftmax	LogSoftmax	LogSoftmax

We offer the model details of the proposed SCLC and compare it with a standard CNN and

Table A.2: Implementation details of UNet (U shape CNN) and SCLC

Layer Name	CNN	SCLC
(Input)	DoubleConv(3,64)	DoubleConv(3,64)
(Down1)	MaxPool2d → DoubleConv(64,128)	Spectral Pool → Spectral DoubleConv(64,128)
(Down2)	MaxPool2d → DoubleConv(128,256)	Spectral Pool → Spectral DoubleConv(128,256)
(Down3)	MaxPool2d → DoubleConv(256,512)	Spectral Pool → Spectral DoubleConv(256,512)
(Down4)	MaxPool2d → DoubleConv(512,1024)	Spectral Pool → Spectral DoubleConv(512,1024)
(Up1)	Upsample→ DoubleConv(1024,512)	Upsample→ DoubleConv(1024,512)
(Up2)	Upsample→ DoubleConv(512,256)	Upsample→ DoubleConv(512,256)
(Up3)	Upsample→ DoubleConv(256,128)	Upsample→ DoubleConv(256,128)
(Up4)	Upsample→ DoubleConv(128,64)	Upsample→ DoubleConv(128,64)
(Output)	Conv2d(64, 1)	Conv2d(64, 1)

Optical fronted CNN in Table.A.1. The models mainly have two parts: 1) convolution block, 2) backend. The convolution block is the three convolution layers with two pooling layer between them. The shared backend is the fully connect layers with a log softmax function. All models has the same backend and the main difference is the computation approach for convolution layers and pooling layers. In the standard CNN, it applies the  $2d$  Convolutional layer and  $2d$  Max-Pooling layer with nonlinear function between layers. While the optical fronted CNN is similar to standard CNN, which only implements the first convolutional layer by the  $4f$  correlator system optically. The SCLC skips the nonlinear functions between layers and computes convolution layers and pooling layer in spectral domain.

Furthermore, we also offer the model details, UNet, in the object classification and show in Table A.2. The Unet network has three parts: 1)Down convolution Layers, 2)Up Convolution layers 3)Skip Connections. A general Down convolution layers is one pooling layer and a double convolutional with batch normalization and nonlinearity between inner layers. When it deploys to SCLC, the pooling and convolution layers will be replaced to the spectral one, and the normalization and nonlinearity will be skipped since it is impossible to be implemented the spectral domain

Table A.3: Accuracy for object classification task tested on STL-10 and Cifar-10

Resolution	STL-10	Cifar-10
$64 \times 64$	19.38%	16.57%
$96 \times 96$	30.1%	58.39%
$128 \times 128$	46.82%	-
$224 \times 224$	64.4%	79.5%

or optical structure. The Up Convolution layer consists of a up sampling layer and two up convolution layers batch normalization and nonlinearity between inner layers. The Skip Connections is between the down block and the up block if they have the same output size. In particular, the first down block will connect to the last up block. Notably, only the Down convolution Layers will be implemented by SCLC or OSCLC. The backend in UNet consists of Up Convolution layers and Skip Connections.

## A.2 STL Task

The STL-10 dataset is also an image classification dataset, which is inspired by the CIFAR-10 dataset but with some modifications. In particular, each class has fewer labeled training examples than in CIFAR-10, but a very large set of unlabeled examples. The primary challenge for STL-10 is for unsupervised image classification or image recognition. The STL-10 has higher resolution of this dataset ( $96 \times 96$ ) than the Cifar-10 ( $32 \times 32$ ). The STL-10 consists of 500 training images (10 pre-defined folds) and 800 test images per class, which is less the Cifar-10 and makes the training harder. We consider the STL-10 dataset to explore the influence of improving the images' resolution. As the larger images can bring more information and hopefully can reach higher performance.

Notably, we explore with Cifar-10 and STL-10 before the HIGH-10 dataset (the high resolution images) to figure out how the model performance improves if the resolution of the input is increased, and show the result in Table.~refstl. As the resolution of STL-10 is  $96 \times 96$ , which is

Table A.4: Subclass accuracy on Cifar-10 and STL-10

task	plane	car	bird	cat	deer	dog	horse	ship	truck	Total accuracy
Cifar-10	84.70%	90.2%	66.9%	57.7%	78.9%	78.9%	82.6%	86.8%	87.4%	80%
STL -10	87.37%	56.75%	78.25%	22.2%	55.4%	20.8%	69%	74.38%	63.0%	64.4%

higher than the Cifar-10 and is expected to contain more information than the Cifar-10. Based on this assumption, we developed this experiment to explore how different input resolution influence the accuracy. For each experiment, the input image resolution is resized to the different resolution from  $64 \times 64$  to  $224 \times 224$ . Because the low resolution images can not directly deploy to the stranded AlexNet (input:  $224 \times 224$ ), the modified AlexNet (less layers) are implemented. We show the result in Table A.3 that increasing the resolution of the input increases the accuracy of final accuracy of the model. Tough promising, we need the explore the high solution (higher than  $224 \times 224$ ) to fully explore this property. Thus, we show this property on HIGH-10 dataset.

In particular, the performance on STL-10 is still low, even increasing the input dimensions. To check the KD performance, we further explore the accuracy on each class and shows in Table.A.4. While the overall performance is higher, each class performance is imbalanced, e.g., the deer class is around 20% lower than other classes. This phenomena also appears in the STL-10 tasks. Based on the observation, this is because the one class is similar to another class. For example, The deer class is similar to house class and the model are not sophisticate enough to predict these two classes. One possible way is to add a focal loss to improve the performance.

### A.3 Object Segmentation Task

We also offer supplementary results for object segmentation tasks in Fig.A.1 and Fig.A.2.




Figure A.1: Examples of object segmentation for the considered benchmarks of Car Segmentation. Left to right: Input image from datasets, Groundtruth, Our work (SCLC/ OSCLC)




Figure A.2: Examples of object segmentation for the considered benchmarks of VOC2012. Top to bottom: Input image from datasets, Groundtruth, Our work (SCLC/ OSCLC)