

CIFAR-10 Image Classification with Deep Learning

Improved Model Architecture and Training Strategy

Anthony Angeles
Santa Clara University

February 3, 2026

a) Model Architecture and Loss Function

CIFAR-10 Data

Let's familiarize ourselves with the data we used from PyTorch's documentation. We will use the CIFAR10 dataset. It has the classes: 'airplane', 'automobile', 'bird', 'cat', 'deer', 'dog', 'frog', 'horse', 'ship', 'truck'. The images in CIFAR-10 are of size 32×32 , i.e. 3-channel color images of 32×32 pixels in size.



Figure 1: Sample Training Data



Figure 2: Ground Truth Images

Model Architecture: ResNet-34 with LeakyReLU Activation

The improved model is based on a **ResNet-34** architecture, modified for the CIFAR-10 dataset. The key architectural choices and improvements are:

1. ResNet-34 Deep Residual Network

The model uses a 34-layer residual network with BasicBlock building blocks. The layer configuration is [3, 4, 6, 3], meaning 3 blocks in layer1, 4 blocks in layer2, 6 blocks in layer3, and 3 blocks in layer4. This results in approximately **21.3 million trainable parameters**.

2. CIFAR-10 Adaptation

Unlike standard ResNet designed for ImageNet (224×224), this implementation is adapted for CIFAR-10's 32×32 images by:

- Using a 3×3 initial convolution with stride=1 instead of a 7×7 convolution with stride=2
- Removing the initial max pooling layer

3. LeakyReLU Activation Function

The model replaces standard ReLU activations with LeakyReLU (negative slope = 0.1). This modification helps prevent the "dying ReLU" problem where neurons can become permanently inactive during training.

Loss Function: Cross-Entropy Loss

The Cross-Entropy Loss function is used for multi-class classification. For a batch of N samples with C classes ($C = 10$ for CIFAR-10)

Training Strategy and Hyperparameters

Data Augmentation

- RandomCrop (32×32 with 4-pixel padding)
- RandomHorizontalFlip ($p = 0.5$)
- RandomErasing/Cutout ($p = 0.5$, scale=0.02–0.33)
- Normalization: $\mu = (0.4914, 0.4822, 0.4465)$,
 $\sigma = (0.2470, 0.2435, 0.2616)$

Optimizer Configuration

- Optimizer: SGD with momentum
- Initial Learning Rate: 0.1122 (via LR Range Test)
- Momentum: 0.9
- Weight Decay: 5×10^{-4} (L2 regularization)
- LR Scheduler: Cosine Annealing ($T_{max} = 200$)

Training Configuration

- Batch Size: 512
- Number of Epochs: 200
- Mixed Precision Training (AMP): Enabled
- Total Training Time: ~ 15 minutes on NVIDIA GeForce RTX 5090

Performance Improvement Summary

Model	Test Accuracy	Parameters	Epochs
Basic CNN (2 conv, 3 fc)	67.35%	$\sim 60K$	100
ResNet-18 with LeakyReLU	86.34%	11.2M	100
ResNet-34 with LeakyReLU	96.00%	21.3M	200

Table 1: Comparison of model architectures and their performance on CIFAR-10

b) Code Implementation

How it works

When we go through the PyTorch documentation code and run it for the first time we hit about 75 percent accuracy, so the first initial steps were to increase batch size and epoch size while still keeping the original CNN that we implemented.

While looking into methods for improving accuracy and learning efficiency, I went ahead and made the model network deeper as well as switching to the ResNet Model, you will see that at first I went with ResNet-18 but later opted for a deeper network, ResNet-34.

Another aspect we changed from the original code is to allow the code itself to run on the GPU and take advantage of the hardware we had access to, an RTX FE 5090 with 32GB of memory.

After we started to change some of the original code to use our own implementations built of PyTorch starting example we saw a drop in accuracy, but that was to be expected as we needed to find the right methods to use in conjunction to reach our desire results, which in my case was a personal goal of at least 95 percent accuracy.

Where we saw the biggest jump in accuracy, and allowing us to reach 96 percent accuracy overall with some classes being 98, was when we switched from ResNet-18 to -34 and when we also introduced a method to optimize the learning rate with Leslie Smith Method. These methods in combination with also adding data augmentation while training, for example flipping the image or clipping it really showed big improvements. These were the biggest changes in my opinion that resulted in such an increase in accuracy, from 67 percent to 96 overall. The code for these demonstrations and the final implementation can be found in a GitHub repository found here: https://github.com/aneangel/improving_CiFAR10.

I must disclose that the methods for tracking multiple test sessions and plot generation, as well as, minor help with brainstorming how to implement the Leslie Smith Method were implemented in conjunction with Claude Sonnet 4.5.

ResNet Model Definition

```
1 class BasicBlockLeakyReLU(nn.Module):
2     expansion = 1
3
4     def __init__(self, in_channels,
5                 out_channels, stride=1,
6                 downsample=None):
7         super().__init__()
8         self.conv1 = nn.Conv2d(
9             in_channels, out_channels,
10            kernel_size=3, stride=stride,
11            padding=1, bias=False)
12        self.bn1 = nn.BatchNorm2d(
13            out_channels)
14        self.leaky_relu = nn.LeakyReLU(
15            negative_slope=0.1, inplace=True)
16        self.conv2 = nn.Conv2d(
17            out_channels, out_channels,
18            kernel_size=3, stride=1,
19            padding=1, bias=False)
20        self.bn2 = nn.BatchNorm2d(
21            out_channels)
22        self.downsample = downsample
23
24    def forward(self, x):
25        identity = x
26        out = self.leaky_relu(
27            self.bn1(self.conv1(x)))
28        out = self.bn2(self.conv2(out))
29        if self.downsample is not None:
30            identity = self.downsample(x)
31        out += identity
32        return self.leaky_relu(out)
```

BasicBlockLeakyReLU Class

```
1 class ResNetLeakyReLU(nn.Module):
2     def __init__(self, block, layers,
3                 num_classes=10):
4         super().__init__()
5         self.in_channels = 64
6         # CIFAR-10 adapted: 3x3 conv
7         self.conv1 = nn.Conv2d(
8             3, 64, kernel_size=3,
9             stride=1, padding=1, bias=False)
10        self.bn1 = nn.BatchNorm2d(64)
11        self.leaky_relu = nn.LeakyReLU(
12            negative_slope=0.1, inplace=True)
13
14        self.layer1 = self._make_layer(
15            block, 64, layers[0], stride=1)
16        self.layer2 = self._make_layer(
17            block, 128, layers[1], stride=2)
18        self.layer3 = self._make_layer(
19            block, 256, layers[2], stride=2)
20        self.layer4 = self._make_layer(
21            block, 512, layers[3], stride=2)
22
23        self.avgpool = nn.AdaptiveAvgPool2d
24        ((1,1))
25        self.fc = nn.Linear(512, num_classes)
26
27        # Kaiming initialization
28        for m in self.modules():
29            if isinstance(m, nn.Conv2d):
30                nn.init.kaiming_normal_(
31                    m.weight, mode='fan_out',
32                    nonlinearity='leaky_relu')
33
34    def resnet34_leaky():
35        return ResNetLeakyReLU(
36            BasicBlockLeakyReLU, [3, 4, 6, 3])
```

ResNetLeakyReLU Class

Data & Training Setup

```
1 # CIFAR-10 statistics
2 cifar10_mean = (0.4914, 0.4822, 0.4465)
3 cifar10_std = (0.2470, 0.2435, 0.2616)
4
5 train_transform = transforms.Compose([
6     transforms.RandomCrop(32, padding=4),
7     transforms.RandomHorizontalFlip(p=0.5),
8     transforms.ToTensor(),
9     transforms.Normalize(
10        cifar10_mean, cifar10_std),
11     transforms.RandomErasing(
12        p=0.5, scale=(0.02, 0.33),
13        ratio=(0.3, 3.3))
14 ])
15
16 test_transform = transforms.Compose([
17     transforms.ToTensor(),
18     transforms.Normalize(
19        cifar10_mean, cifar10_std)
20 ])
```

Data Augmentation Pipeline

```
1 # Hyperparameters
2 learning_rate = 0.1122 # LR Range Test
3 momentum = 0.9
4 num_epochs = 200
5 weight_decay = 5e-4
6 batch_size = 512
7
8 # Model, Loss, Optimizer
9 net = resnet34_leaky().to(device)
10 criterion = nn.CrossEntropyLoss()
11 optimizer = optim.SGD(
12     net.parameters(), lr=learning_rate,
13     momentum=momentum,
14     weight_decay=weight_decay)
15 scheduler = optim.lr_scheduler.
16     CosineAnnealingLR(
17         optimizer, T_max=num_epochs, eta_min=1e
18         -4)
19
20 # Mixed Precision Training
21 scaler = torch.amp.GradScaler(
22     'cuda', enabled=True)
```

Training Configuration

```
1 for epoch in range(num_epochs):
2     for inputs, labels in trainloader:
3         inputs = inputs.to(device)
4         labels = labels.to(device)
5         optimizer.zero_grad()
6
7         # AMP forward pass
8         with torch.amp.autocast(
9             'cuda', enabled=True):
10            outputs = net(inputs)
11            loss = criterion(outputs, labels)
12
13         # Scaled backward pass
14         scaler.scale(loss).backward()
15         scaler.step(optimizer)
16         scaler.update()
17
18     scheduler.step()
```

Training Loop with AMP

Evaluation Code (Overall + Per-Class Accuracy)

```

1 # Initialize counters
2 correct = 0
3 total = 0
4 class_correct = {classname: 0 for classname in classes}
5 class_total = {classname: 0 for classname in classes}
6
7 with torch.no_grad():
8     for images, labels in testloader:
9         images, labels = images.to(device), labels.to(device)
10        outputs = net(images)
11        _, predicted = torch.max(outputs.data, 1)
12        total += labels.size(0)
13        correct += (predicted == labels).sum().item()
14
15        # Per-class accuracy
16        for label, prediction in zip(labels, predicted):
17            classname = classes[label]
18            class_total[classname] += 1
19            if label == prediction:
20                class_correct[classname] += 1
21
22 test_accuracy = 100 * correct / total
23 print(f'Overall Accuracy: {test_accuracy:.2f}%')
24
25 # Display per-class accuracy
26 for classname in classes:
27     acc = 100 * class_correct[classname] / class_total[classname]
28     print(f'{classname}: {class_correct[classname]}/{class_total[classname]} = {acc:.2f}%')

```

c) Screenshots and Results

Overall Network Accuracy

The improved ResNet-34 with LeakyReLU achieved a test accuracy of **96.00%** on the CIFAR-10 dataset, correctly classifying 9,600 out of 10,000 test images. This represents a significant improvement from the baseline Basic CNN (67.35%) and the intermediate ResNet-18 model (86.34%).

Per-Class Accuracy Results

The per-class accuracy results demonstrate strong performance across all 10 CIFAR-10 classes, with accuracy ranging from 90.90% (cat) to 98.40% (car). These are the actual results from running the code:

Class	Correct / Total	Accuracy
plane	964 / 1000	96.40%
car	984 / 1000	98.40%
bird	950 / 1000	95.00%
<i>cat</i>	<i>909 / 1000</i>	<i>90.90%</i>
deer	970 / 1000	97.00%
dog	932 / 1000	93.20%
frog	974 / 1000	97.40%
horse	974 / 1000	97.40%
ship	972 / 1000	97.20%
truck	971 / 1000	97.10%




Figure 3: Per Class Accuracies

Table 2: Per-class accuracy. **Bold**: Best (car).

Italic: Worst (cat).

Best Class: car (98.40%)

Worst Class: cat (90.90%)

Experiment Tracking Summary

Exp #	Timestamp	Model	Accuracy	LR	Epochs
2	2026-02-02 11:45	ResNet-34-LeakyReLU	96.00%	0.1122	200
1	2026-02-02 10:22	ResNet-18-LeakyReLU	86.34%	0.0100	100
0	2026-02-02 09:59	Basic CNN	67.35%	0.0010	100

Table 3: Experiment tracking showing progressive improvement in model accuracy