

CPEN 455: Deep Learning

Tutorial 9: Conditional PixelCNN++

Felix Fu

University of British Columbia

Winter, Term 2, 2025

Outline

- Intro to Conda environment
- **Go through important functions**
- **How to integrate labels and picture?**
- Potential numerical issue
- Important Hyperparameters
- **How long it takes to train a conditional PixelCNN++?**
- **Interfaces + Takeaways**
- **Questions?**

Intro to Conda Environment

- [Install Pytorch](#)
- `pip install -r requirements.txt`

INSTALL PYTORCH

Select your preferences and run the install command. Stable represents the most currently tested and supported version of PyTorch. This should be suitable for many users. Preview is available if you want the latest, not fully tested and supported, builds that are generated nightly. Please ensure that you have **met the prerequisites below (e.g., numpy)**, depending on your package manager. Anaconda is our recommended package manager since it installs all dependencies. You can also [install previous versions of PyTorch](#). Note that LibTorch is only available for C++.

NOTE: Latest PyTorch requires Python 3.8 or later. For more details, see Python section below.

PyTorch Build	Stable (2.2.2)		Preview (Nightly)	
Your OS	Linux	Mac	Windows	
Package	Conda	Pip	LibTorch	Source
Language	Python		C++ / Java	
Compute Platform	CUDA 11.8	CUDA 12.1	ROCm 5.7	Default
Run this Command:	<code>pip3 install torch torchvision torchaudio</code>			

Intro to Conda Environment

- [Install Pytorch](#)
- `pip install -r requirements.txt`

INSTALL PYTORCH

Select your preferences and run the install command. Stable represents the most currently tested and supported version of PyTorch. This should be suitable for many users. Preview is available if you want the latest, not fully tested and supported, builds that are generated nightly. Please ensure that you have **met the prerequisites below (e.g., numpy)**, depending on your package manager. Anaconda is our recommended package manager since it installs all dependencies. You can also [install previous versions of PyTorch](#). Note that LibTorch is only available for C++.

NOTE: Latest PyTorch requires Python 3.8 or later. For more details, see Python section below.

PyTorch Build	<input checked="" type="checkbox"/> Stable (2.2.2)	<input type="checkbox"/> Preview (Nightly)		
Your OS	<input type="checkbox"/> Linux	<input type="checkbox"/> Mac	<input checked="" type="checkbox"/> Windows	
Package	<input type="checkbox"/> Conda	<input checked="" type="checkbox"/> Pip	<input type="checkbox"/> LibTorch	<input type="checkbox"/> Source
Language	<input checked="" type="checkbox"/> Python		<input type="checkbox"/> C++ / Java	
Compute Platform	<input type="checkbox"/> CUDA 11.8	<input checked="" type="checkbox"/> CUDA 12.1	<input type="checkbox"/> ROCm 5.7	<input type="checkbox"/> CPU
Run this Command:	<code>pip3 install torch torchvision torchaudio --index-url https://download.pytorch.org/whl/cu121</code>			

[Previous versions of PyTorch](#) >

nvidia-smi

```
qihang — qihangz@ece-cdl-lw02:~ — ssh works — 81x26
(base) → ~ nvidia-smi
Mon Apr 8 10:39:30 2024
+-----+-----+-----+-----+-----+-----+-----+
| NVIDIA-SMI 520.61.05      | Driver Version: 520.61.05 | CUDA Version: 11.8 |
+-----+-----+-----+-----+-----+-----+-----+
| GPU   Name                   Persistence-M | Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap |      Memory-Usage | GPU-Util  Compute M. |
|====+=====+====+=====+=====+=====+=====+
|  0   NVIDIA GeForce ...     On          | 00000000:01:00.0 Off  |           N/A       |
|  0%   32C    P8      23W / 350W | 362MiB / 24576MiB |      0%      Default |
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+
| Processes: |
| GPU   GI    CI          PID    Type   Process name                      GPU Memory |
|====+=====+====+=====+=====+=====+=====+
|  0     N/A  N/A         2228     G   /usr/lib/xorg/Xorg                 35MiB |
|  0     N/A  N/A         6946     G   /usr/lib/xorg/Xorg                 190MiB |
|  0     N/A  N/A         7118     G   /usr/bin/gnome-shell                16MiB |
|  0     N/A  N/A        64601     G   ...on=20240201-180133.047000       10MiB |
|  0     N/A  N/A        64670     G   ...RendererForSitePerProcess        8MiB |
+-----+-----+-----+-----+-----+-----+-----+
(base) → ~
```

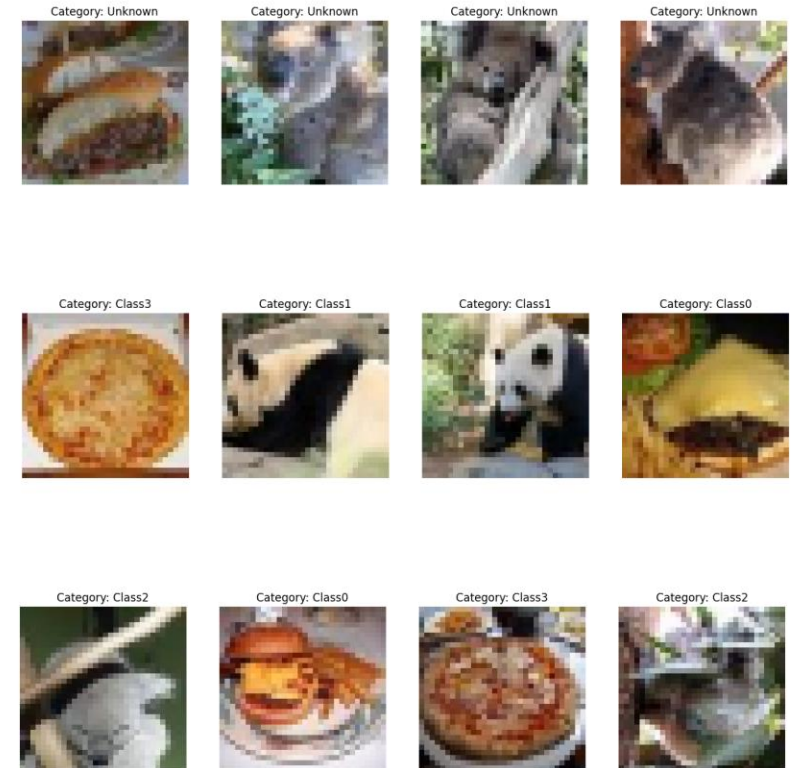
Dataset.py

```
dataset.py x
dataset.py > ...
21 class CPEN455Dataset(Dataset):
22     def __init__(self, root_dir, mode='train', transform=None):
23
24         Args:
25             root_dir (string): Directory with all the images and labels.
26             transform (callable, optional): Optional transform to be applied on a sample.
27         """
28         ROOT_DIR = './data'
29         root_dir = os.path.join(root_dir, mode)
30         self.root_dir = root_dir
31         self.transform = transform
32         self.samples = [] # List to store image paths along with domain and category
33         # Walk through the directory structure
34         csv_path = os.path.join(ROOT_DIR, mode + '.csv')
35         df = pd.read_csv(csv_path, header=None, names=['path', 'label'])
36         # Convert DataFrame to a list of tuples
37         self.samples = list(df.itertuples(index=False, name=None))
38         self.samples = [(os.path.join(ROOT_DIR, path), label) for path, label in self.samples]
39
40     def __len__(self):
41         return len(self.samples)
42
43     def __getitem__(self, idx):
44         img_path, category = self.samples[idx]
45         if category in my_bidict.values():
46             category_name = my_bidict.inverse[category]
47         else:
48             category_name = "Unknown"
49         # print(img_path)
50         image = read_image(img_path) # Reads the image as a tensor
51         image = image.type(torch.float32) / 255. # Normalize to [0, 1]
52         if image.shape[0] == 1:
53             image = replicate_color_channel(image)
54         if self.transform:
55             image = self.transform(image)
56         return image, category_name
```

Dataset.py

```
dataset.py M x
dataset.py > show_images

68
69 if __name__ == '__main__':
70
71     transform_32 = Compose([
72         Resize((32, 32)), # Resize images to 32 * 32
73         rescaling
74     ])
75     dataset_list = ['train', 'validation', 'test']
76
77     for mode in dataset_list:
78         print(f"Mode: {mode}")
79         dataset = CPEN455Dataset(root_dir='./data', transform=transform_32, mode=mode)
80         data_loader = DataLoader(dataset, batch_size = 4, shuffle=True)
81         # Sample from the DataLoader
82         for images, categories in tqdm(data_loader):
83             print(images.shape, categories)
84             images = torch.round(rescaling_inv(images) * 255).type(torch.uint8)
85             show_images(images, categories, mode)
86             break # We only want to see one batch of 4 images in this example
87
```

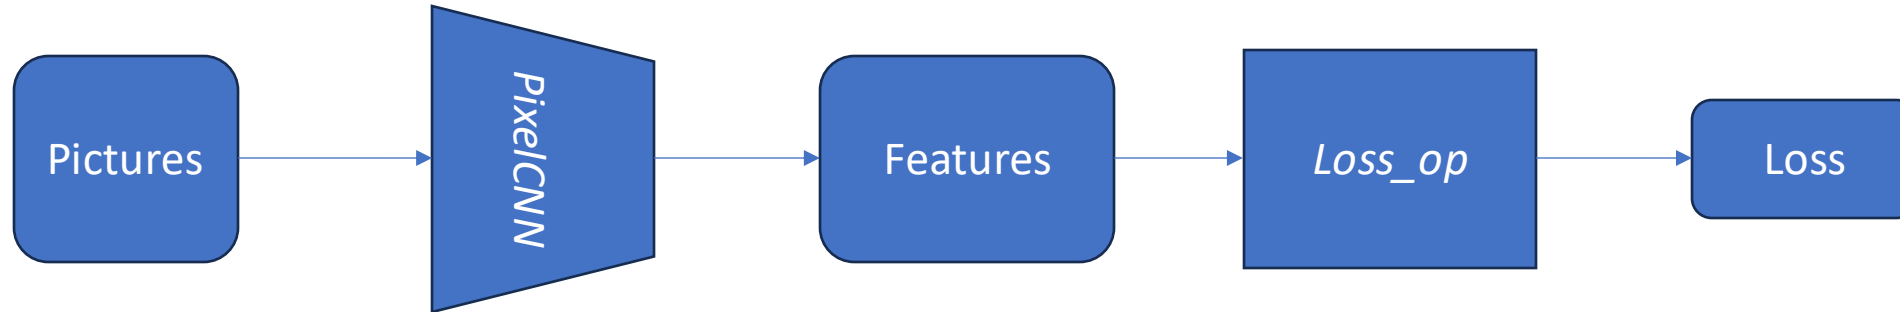


pcnn_train.py

```
dataset.py pcnn_train.py x utils.py model.py
pcnn_train.py > ...
16
17 def train_or_test(model, data_loader, optimizer, loss_op, device, args, epoch, mode = 'training'):
18     if mode == 'training':
19         model.train()
20     else:
21         model.eval()
22
23     deno = args.batch_size * np.prod(args.obs) * np.log(2.)
24     loss_tracker = mean_tracker()
25
26     for batch_idx, item in enumerate(tqdm(data_loader)):
27         model_input, _ = item → label
28         model_input = model_input.to(device)
29         model_output = model(model_input) ←
30         loss = loss_op(model_input, model_output)
31         loss_tracker.update(loss.item()/deno)
32         if mode == 'training':
33             optimizer.zero_grad()
34             loss.backward()
35             optimizer.step()
36
37     if args.en_wandb:
38         wandb.log({mode + "-Average-BPD" : loss_tracker.get_mean()})
39         wandb.log({mode + "-epoch": epoch})
```

Most important function in whole project:

discretized_mix_logistic_loss(x, l) line36 in utils.py



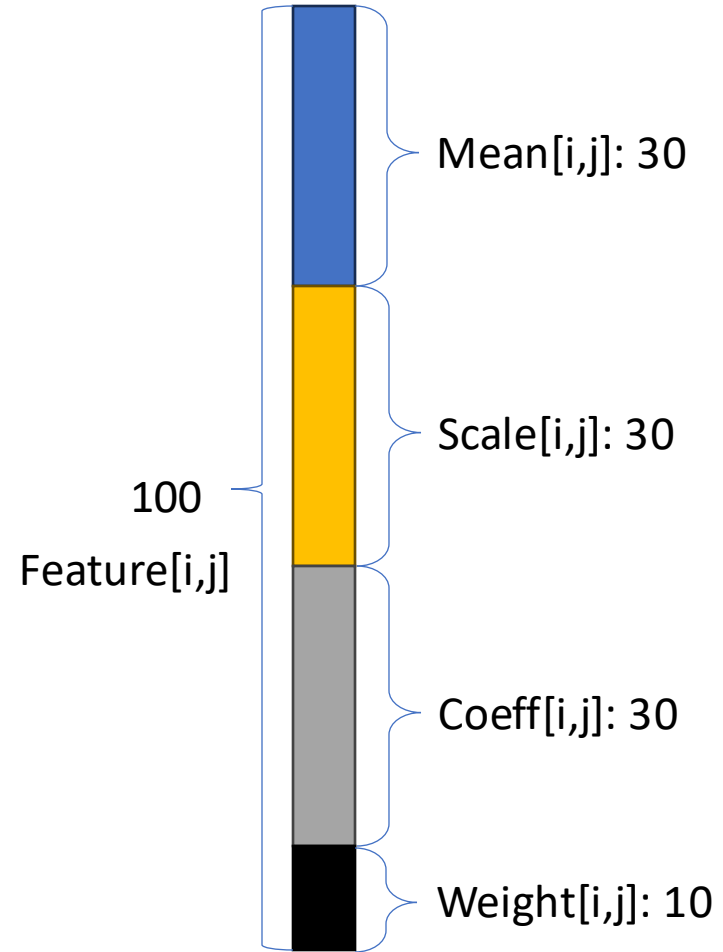
$$\begin{aligned} D_{KL}(P_{data}, P_{\theta}) &= \sum_{\mathcal{Z}} P_{data}(\mathcal{Z}) \cdot \log \frac{P_{data}(\mathcal{Z})}{P_{\theta}(\mathcal{Z})} \\ &= - \sum_{\mathcal{Z}} P_{data}(\mathcal{Z}) \log P_{\theta}(\mathcal{Z}) + \underbrace{\sum_{\mathcal{Z}} P_{data}(\mathcal{Z}) \log P_{data}(\mathcal{Z})}_{\text{constant}} \\ &\approx - \sum_{i=1}^N \log P_{\theta}(i^{\text{th}} \text{ picture}) \end{aligned}$$

Short conclusion: the output of loss function should be of the form NLL

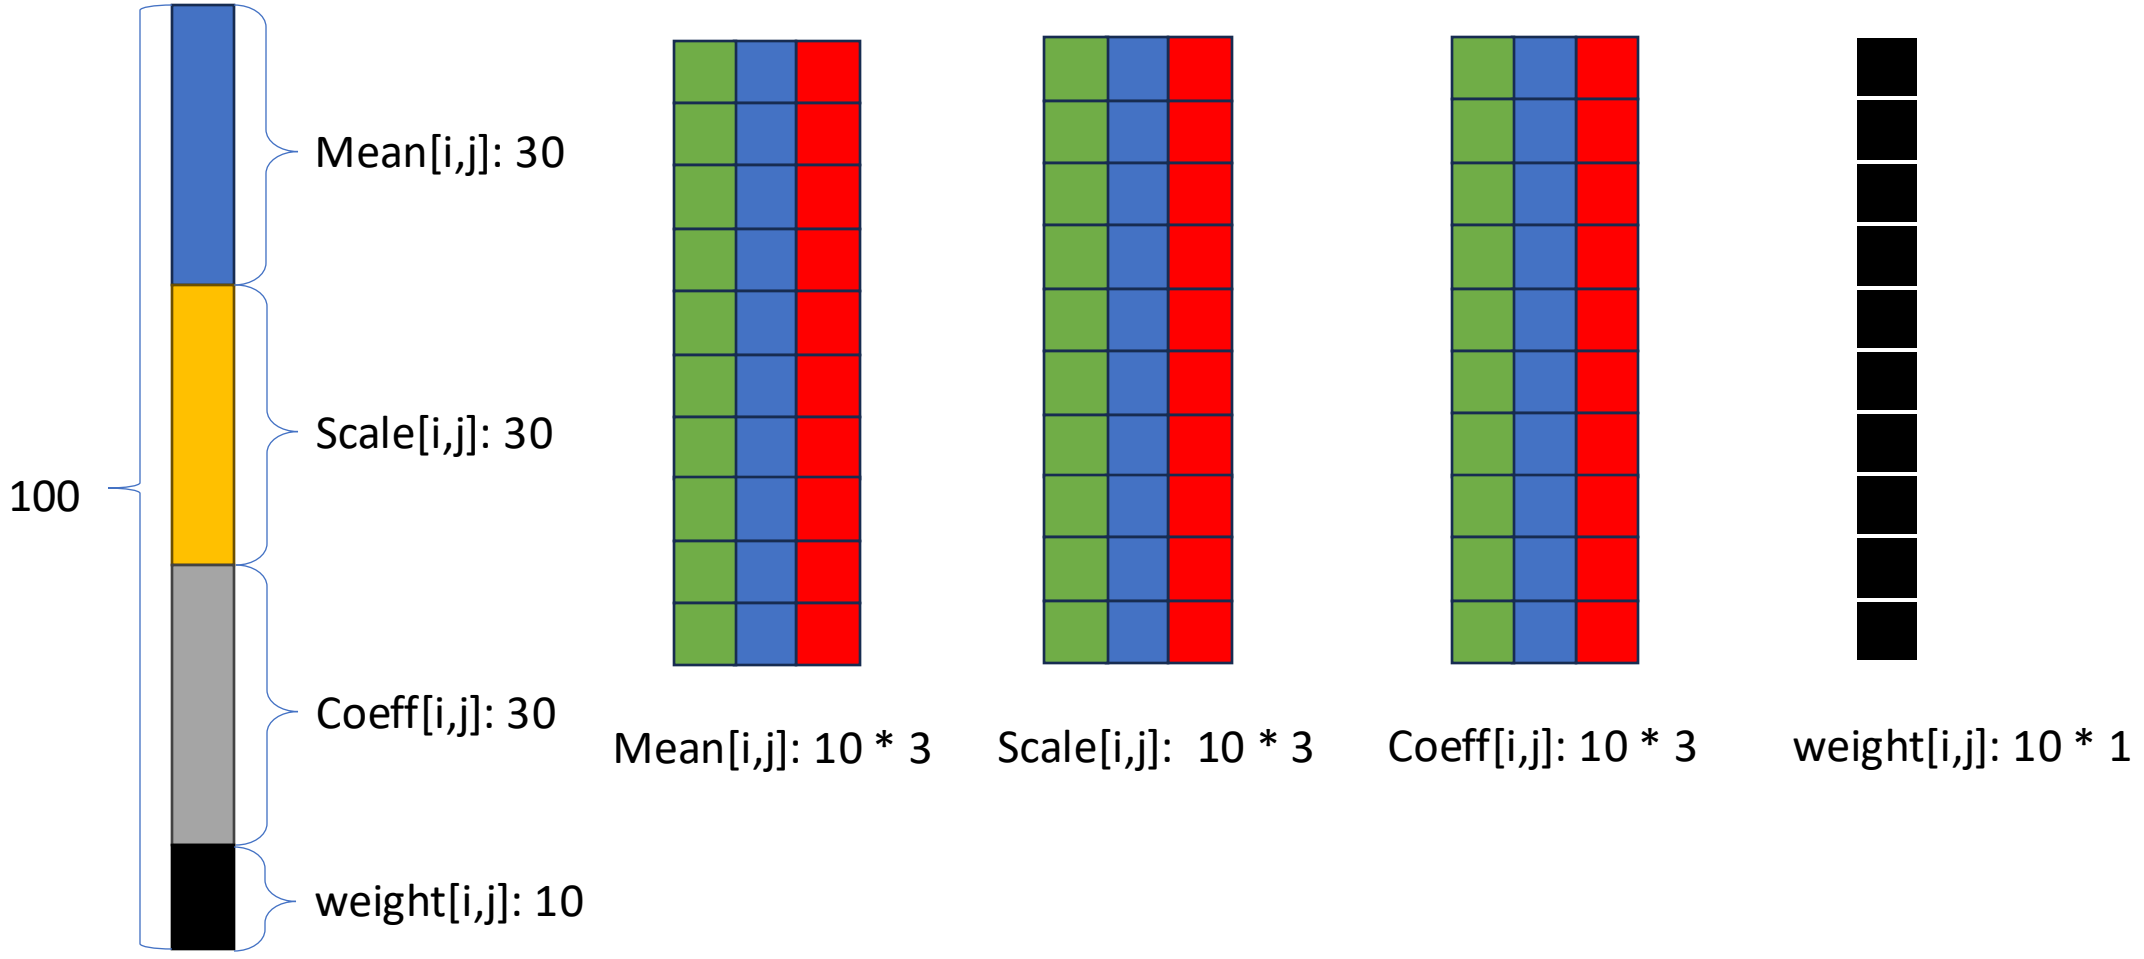
discretized_mix_logistic_loss(x, l)

Feature
(output of
PCNN)

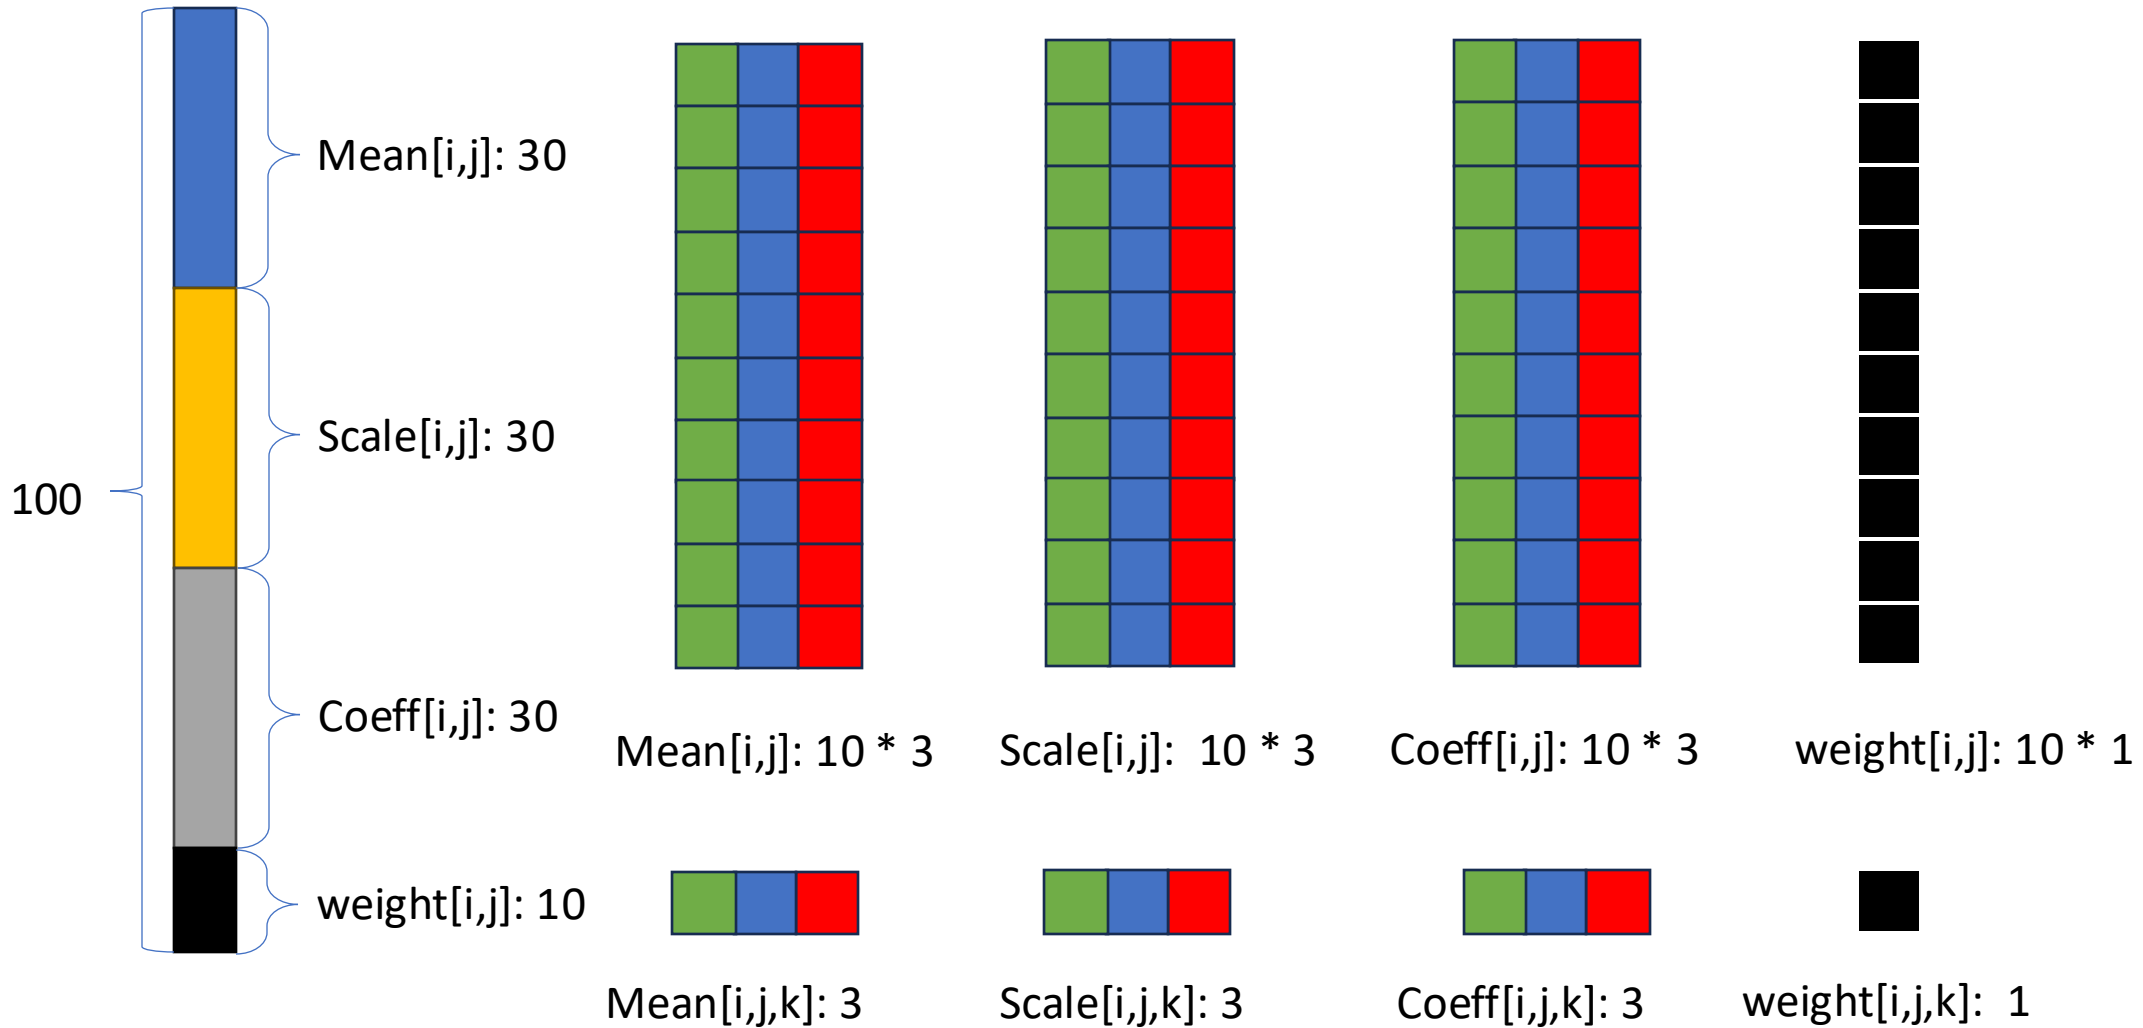
H * W * 100



discretized_mix_logistic_loss(x, l)



discretized_mix_logistic_loss(x, l)



discretized_mix_logistic_loss(x, l)



Mean[i,j,k]: 3



Scale[i,j,k]: 3



Coeff[i,j,k]: 3



weight[i,j,k]: 1

$$p(r_{i,j}, g_{i,j}, b_{i,j} | C_{i,j}) = P(r_{i,j} | \mu_r(C_{i,j}), s_r(C_{i,j})) \times P(g_{i,j} | \mu_g(C_{i,j}, r_{i,j}), s_g(C_{i,j})) \\ \times P(b_{i,j} | \mu_b(C_{i,j}, r_{i,j}, g_{i,j}), s_b(C_{i,j}))$$
$$\mu_g(C_{i,j}, r_{i,j}) = \mu_g(C_{i,j}) + \alpha(C_{i,j})r_{i,j}$$
$$\mu_b(C_{i,j}, r_{i,j}, g_{i,j}) = \mu_b(C_{i,j}) + \beta(C_{i,j})r_{i,j} + \gamma(C_{i,j})g_{i,j}$$

discretized_mix_logistic_loss(x, l)



Mean[i,j,k]: 3



Scale[i,j,k]: 3



Coeff[i,j,k]: 3



weight[i,j,k]: 1

Mean[i,j,k,v], v = 0,1,2

Scale[i,j,k,v], v = 0,1,2

Coeff[i,j,k,v], v = 0,1,2



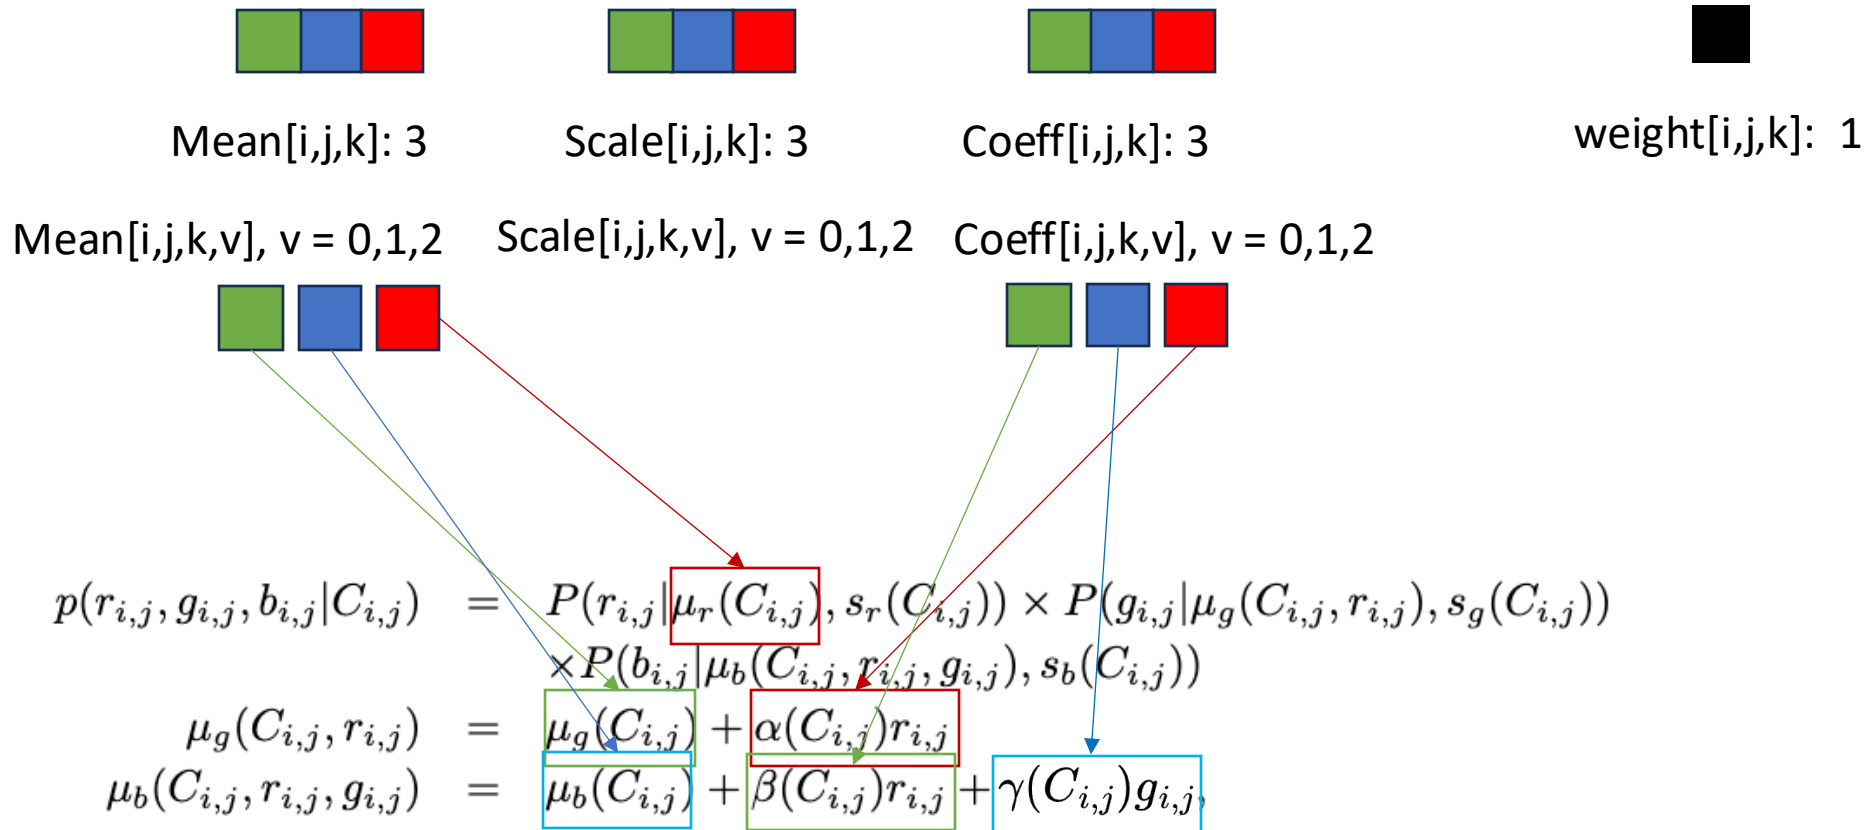
$$p(r_{i,j}, g_{i,j}, b_{i,j} | C_{i,j}) = P(r_{i,j} | \mu_r(C_{i,j}), s_r(C_{i,j})) \times P(g_{i,j} | \mu_g(C_{i,j}, r_{i,j}), s_g(C_{i,j})) \\ \times P(b_{i,j} | \mu_b(C_{i,j}, r_{i,j}, g_{i,j}), s_b(C_{i,j}))$$

$$\mu_g(C_{i,j}, r_{i,j}) = \mu_g(C_{i,j}) + \alpha(C_{i,j})r_{i,j}$$

$$\mu_b(C_{i,j}, r_{i,j}, g_{i,j}) = \mu_b(C_{i,j}) + \beta(C_{i,j})r_{i,j} + \gamma(C_{i,j})g_{i,j}$$

Note: equation from the original paper has a typo.

discretized_mix_logistic_loss(x, l)



discretized_mix_logistic_loss(x, l)

Mean[i,j,k,v], v = 0,1,2

Scale[i,j,k,v], v = 0,1,2

Coeff[i,j,k,v], v = 0,1,2



weight[i,j,k]: 1

$$p(r_{i,j}, g_{i,j}, b_{i,j} | C_{i,j}) = P(r_{i,j} | \mu_r(C_{i,j}), s_r(C_{i,j})) \times P(g_{i,j} | \mu_g(C_{i,j}, r_{i,j}), s_g(C_{i,j})) \\ \times P(b_{i,j} | \mu_b(C_{i,j}, r_{i,j}, g_{i,j}), s_b(C_{i,j}))$$

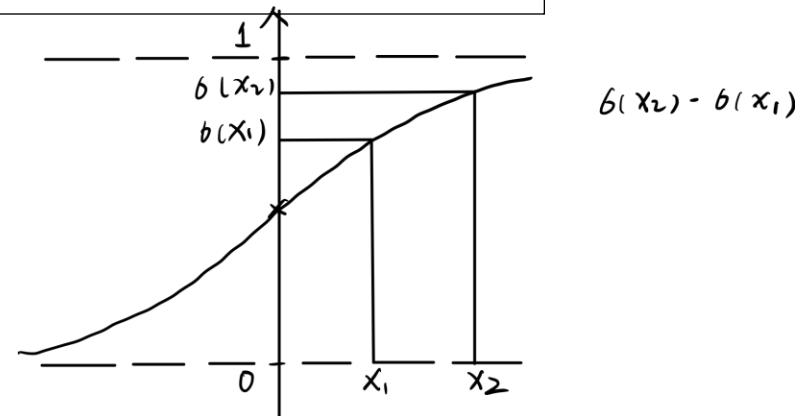
$$\mu_g(C_{i,j}, r_{i,j}) = \mu_g(C_{i,j}) + \alpha(C_{i,j})r_{i,j}$$

$$\mu_b(C_{i,j}, r_{i,j}, g_{i,j}) = \mu_b(C_{i,j}) + \beta(C_{i,j})r_{i,j} + \gamma(C_{i,j})g_{i,j}$$

$$P(x | \pi, \mu, s) = \sum_{i=1}^K \pi_i \left[\sigma\left(\frac{x + 0.5 - \mu_i}{s_i}\right) - \sigma\left(\frac{x - 0.5 - \mu_i}{s_i}\right) \right]$$

Scale

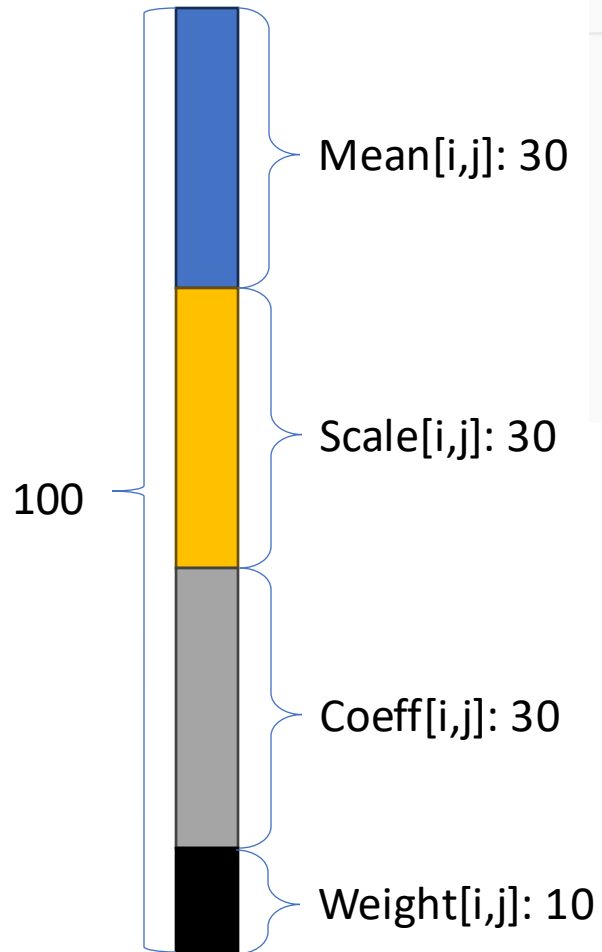
Mean



discretized_mix_logistic_loss(x, l)

Feature
(output of
PCNN)

H * W * 100



```
36 def discretized_mix_logistic_loss(x, l):  
43  
44     # here and below: unpacking the params of the mixture of logistics  
45     nr_mix = int(ls[-1] / 10)  
46     logit_probs = l[:, :, :, :nr_mix]  
47     l = l[:, :, :, nr_mix:].contiguous().view(xs + [nr_mix * 3]) # 3 for mean, scale, coef  
48     means = l[:, :, :, :nr_mix]  
49     # log_scales = torch.max(l[:, :, :, :, nr_mix:2 * nr_mix], -7.)  
50     log_scales = torch.clamp(l[:, :, :, :, nr_mix:2 * nr_mix], min=-7.)  
51  
52     coeffs = F.tanh(l[:, :, :, :, 2 * nr_mix:3 * nr_mix])
```


discretized_mix_logistic_loss(x, l)

Mean[i,j,k,v], v = 0,1,2



Coeff[i,j,k,v], v = 0,1,2



$$p(r_{i,j}, g_{i,j}, b_{i,j} | C_{i,j}) = P(r_{i,j} | \mu_r(C_{i,j}), s_r(C_{i,j})) \times P(g_{i,j} | \mu_g(C_{i,j}, r_{i,j}), s_g(C_{i,j})) \\ \times P(b_{i,j} | \mu_b(C_{i,j}, r_{i,j}, g_{i,j}), s_b(C_{i,j}))$$

$$\mu_g(C_{i,j}, r_{i,j}) = \mu_g(C_{i,j}) + \alpha(C_{i,j})r_{i,j}$$

$$\mu_b(C_{i,j}, r_{i,j}, g_{i,j}) = \mu_b(C_{i,j}) + \beta(C_{i,j})r_{i,j} + \gamma(C_{i,j})g_{i,j}$$

```
36 def discretized_mix_logistic_loss(x, l):
```

```
51
```

```
52     coeffs = F.tanh(l[:, :, :, :, 2 * nr_mix:3 * nr_mix])
```

```
53     # here and below: getting the means and adjusting them based on preceding
```

```
54     # sub-pixels
```

```
55     x = x.contiguous() You, 2 weeks ago • first commit
```

```
56     x = x.unsqueeze(-1) + Variable(torch.zeros(xs + [nr_mix]).to(x.device), requires_grad=False)
```

```
57     m2 = (means[:, :, :, 1, :] + coeffs[:, :, :, 0, :])
```

```
58     |   |   |   * x[:, :, :, 0, :]).view(xs[0], xs[1], xs[2], 1, nr_mix)
```

```
59
```

```
60     m3 = (means[:, :, :, 2, :] + coeffs[:, :, :, 1, :] * x[:, :, :, 0, :] +
```

```
61     |   |   |   coeffs[:, :, :, 2, :] * x[:, :, :, 1, :]).view(xs[0], xs[1], xs[2], 1, nr_mix)
```

```
62
```

```
63     means = torch.cat((means[:, :, :, 0, :].unsqueeze(3), m2, m3), dim=3)
```

discretized_mix_logistic_loss(x, l)

```
36 def discretized_mix_logistic_loss(x, l):
64     centered_x = x - means
65     inv_stdv = torch.exp(-log_scales)
66     plus_in = inv_stdv * (centered_x + 1. / 255.)
67     cdf_plus = F.sigmoid(plus_in)
68     min_in = inv_stdv * (centered_x - 1. / 255.)
69     cdf_min = F.sigmoid(min_in)
70     # log probability for edge case of 0 (before scaling)
71     log_cdf_plus = plus_in - F.softplus(plus_in)
72     # log probability for edge case of 255 (before scaling)
73     log_one_minus_cdf_min = -F.softplus(min_in)
74     cdf_delta = cdf_plus - cdf_min # probability for all other cases
```

$$[\sigma((x + 0.5 - \mu_i)/s_i) - \sigma((x - 0.5 - \mu_i)/s_i)]$$

discretized_mix_logistic_loss(x, l)

```
36 def discretized_mix_logistic_loss(x, l):
75     mid_in = inv_stdv * centered_x
76     # log probability in the center of the bin, to be used in extreme c
77     # (not actually used in our code)
78     log_pdf_mid = mid_in - log_scales - 2. * F.softplus(mid_in)
79
80     # now select the right output: left edge case, right edge case, nor
81     # case, extremely low prob case (doesn't actually happen for us)
82
83     # this is what we are really doing, but using the robust version be
84     # log_probs = tf.select(x < -0.999, log_cdf_plus, tf.select(x > 0.9
85
86     # robust version, that still works if probabilities are below 1e-5
87     # tensorflow backpropagates through tf.select() by multiplying with
88     # the 1e-12 in tf.maximum(cdf_delta, 1e-12) is never actually used
89     # if the probability on a sub-pixel is below 1e-5, we use an approx
90     # based on the assumption that the log-density is constant in the l
91     # the observed sub-pixel value
92
93     inner_inner_cond = (cdf_delta > 1e-5).float()
94     inner_inner_out = inner_inner_cond * torch.log(torch.clamp(cdf_de
95     inner_cond = (x > 0.999).float()
96     inner_out = inner_cond * log_one_minus_cdf_min + (1. - inner
97     cond = (x < -0.999).float()
98     log_probs = cond * log_cdf_plus + (1. - cond) * inner_out
```

For numerical stability

discretized_mix_logistic_loss(x, l)

```
def discretized_mix_logistic_loss(x, l):  
    log_probs = torch.sum(log_probs, dim=3) + log_prob_from_logits(logit_probs)  
  
    return -torch.sum(log_sum_exp(log_probs))
```

$$P(x|\pi, \mu, s) = \sum_{i=1}^K \pi_i [\sigma((x + 0.5 - \mu_i)/s_i) - \sigma((x - 0.5 - \mu_i)/s_i)]$$

$$p(r_{i,j}, g_{i,j}, b_{i,j}|C_{i,j}) = P(r_{i,j}|\mu_r(C_{i,j}), s_r(C_{i,j})) \times P(g_{i,j}|\mu_g(C_{i,j}, r_{i,j}), s_g(C_{i,j})) \\ \times P(b_{i,j}|\mu_b(C_{i,j}, r_{i,j}, g_{i,j}), s_b(C_{i,j}))$$

Short conclusion:

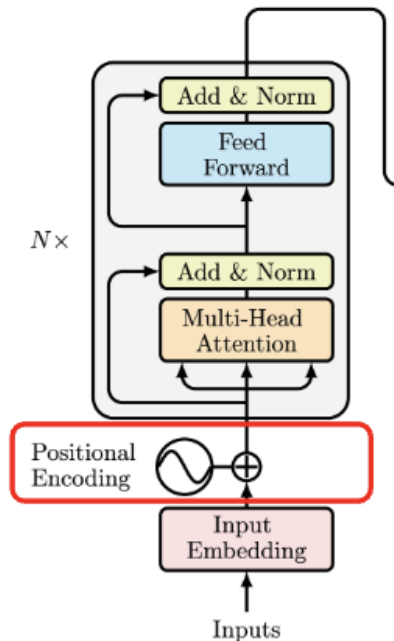
If you want to get the log prob of each picture, you should modify the last line

How to integrate labels and picture?

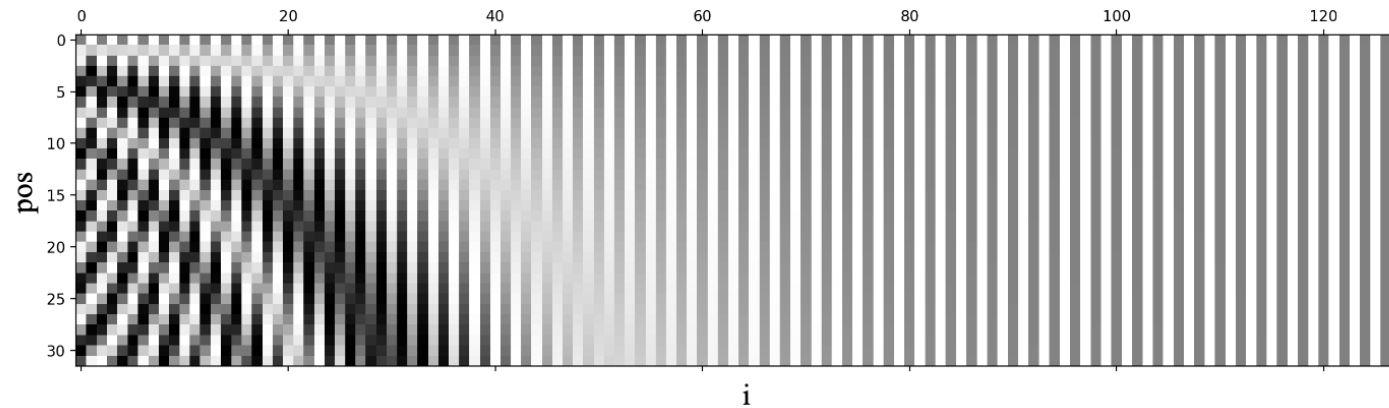
- How to encode labels?
- How to integrate it with pictures?

some hints from Positional Encoding:

1. Encode the label into a tensor (positional encoding in PA2)
2. Integrate it with input embedding



Positional Encoding

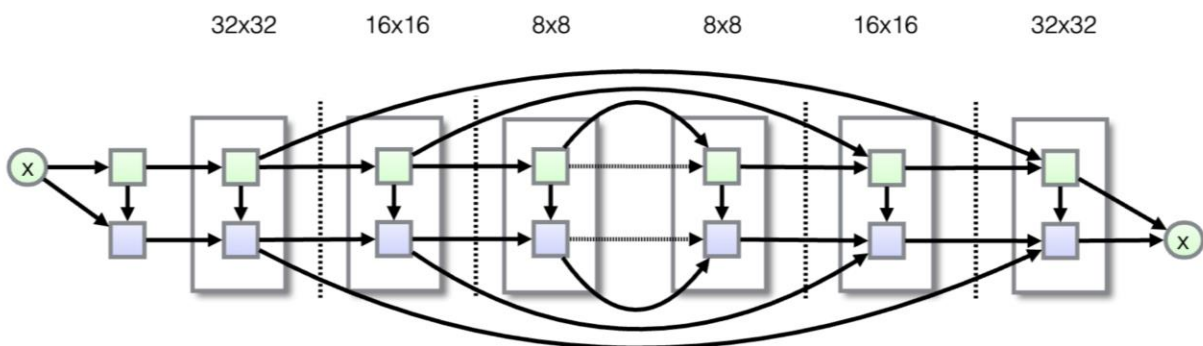


$$PE_{(pos,2i)} = \sin(pos/10000^{2i/d_{model}})$$

$$PE_{(pos,2i+1)} = \cos(pos/10000^{2i/d_{model}})$$

Text and image integration

Different position has different impact
explore it by yourself.



```
53 class PixelCNN(nn.Module):
100     def forward(self, x, sample=False):
112
113         ### UP PASS ###
114         x = x if sample else torch.cat((x, self.init_padding), 1)
115         u_list = [self.u_init(x)]
116         ul_list = [self.ul_init[0](x) + self.ul_init[1](x)]
117         for i in range(3):
118             # resnet block
119             u_out, ul_out = self.up_layers[i](u_list[-1], ul_list[-1])
120             u_list += u_out
121             ul_list += ul_out
122
123             if i != 2:
124                 # downscale (only twice)
125                 u_list += [self.downsize_u_stream[i](u_list[-1])]
126                 ul_list += [self.downsize_ul_stream[i](ul_list[-1])]
127
128         ### DOWN PASS ###
129         u = u_list.pop()
130         ul = ul_list.pop()
131
132         for i in range(3):
133             # resnet block
134             u, ul = self.down_layers[i](u, ul, u_list, ul_list)
135
136             # upscale (only twice)
137             if i != 2 :
138                 u = self.upscale_u_stream[i](u)
139                 ul = self.upscale_ul_stream[i](ul)
140
141         x_out = self.nin_out(F.elu(ul))
```

Potential numerical issue

$$P(i|x) = \frac{P(x|i)}{\sum_{k=1}^4 P(x|k)}$$

$$P(x|i) = \exp(\log P(x|i))$$

$$\log P(x|i)$$

$$= \sum \log P(x_j | x_{<j}, i)$$

$$\text{if } \log P(x_j | x_{<j}, i) \approx -8$$

$$\log P(x|i) \approx -8 \times \underbrace{32 \times 32 \times 32}_{\text{}} = -24k$$

$$\exp(-24k) \rightarrow \text{underflow}$$

$$\log P(i|x) = \log P(x|i) - \underbrace{\log \sum_{k=1}^4 P(x|k)}_{\text{torch.logsumexp}}$$

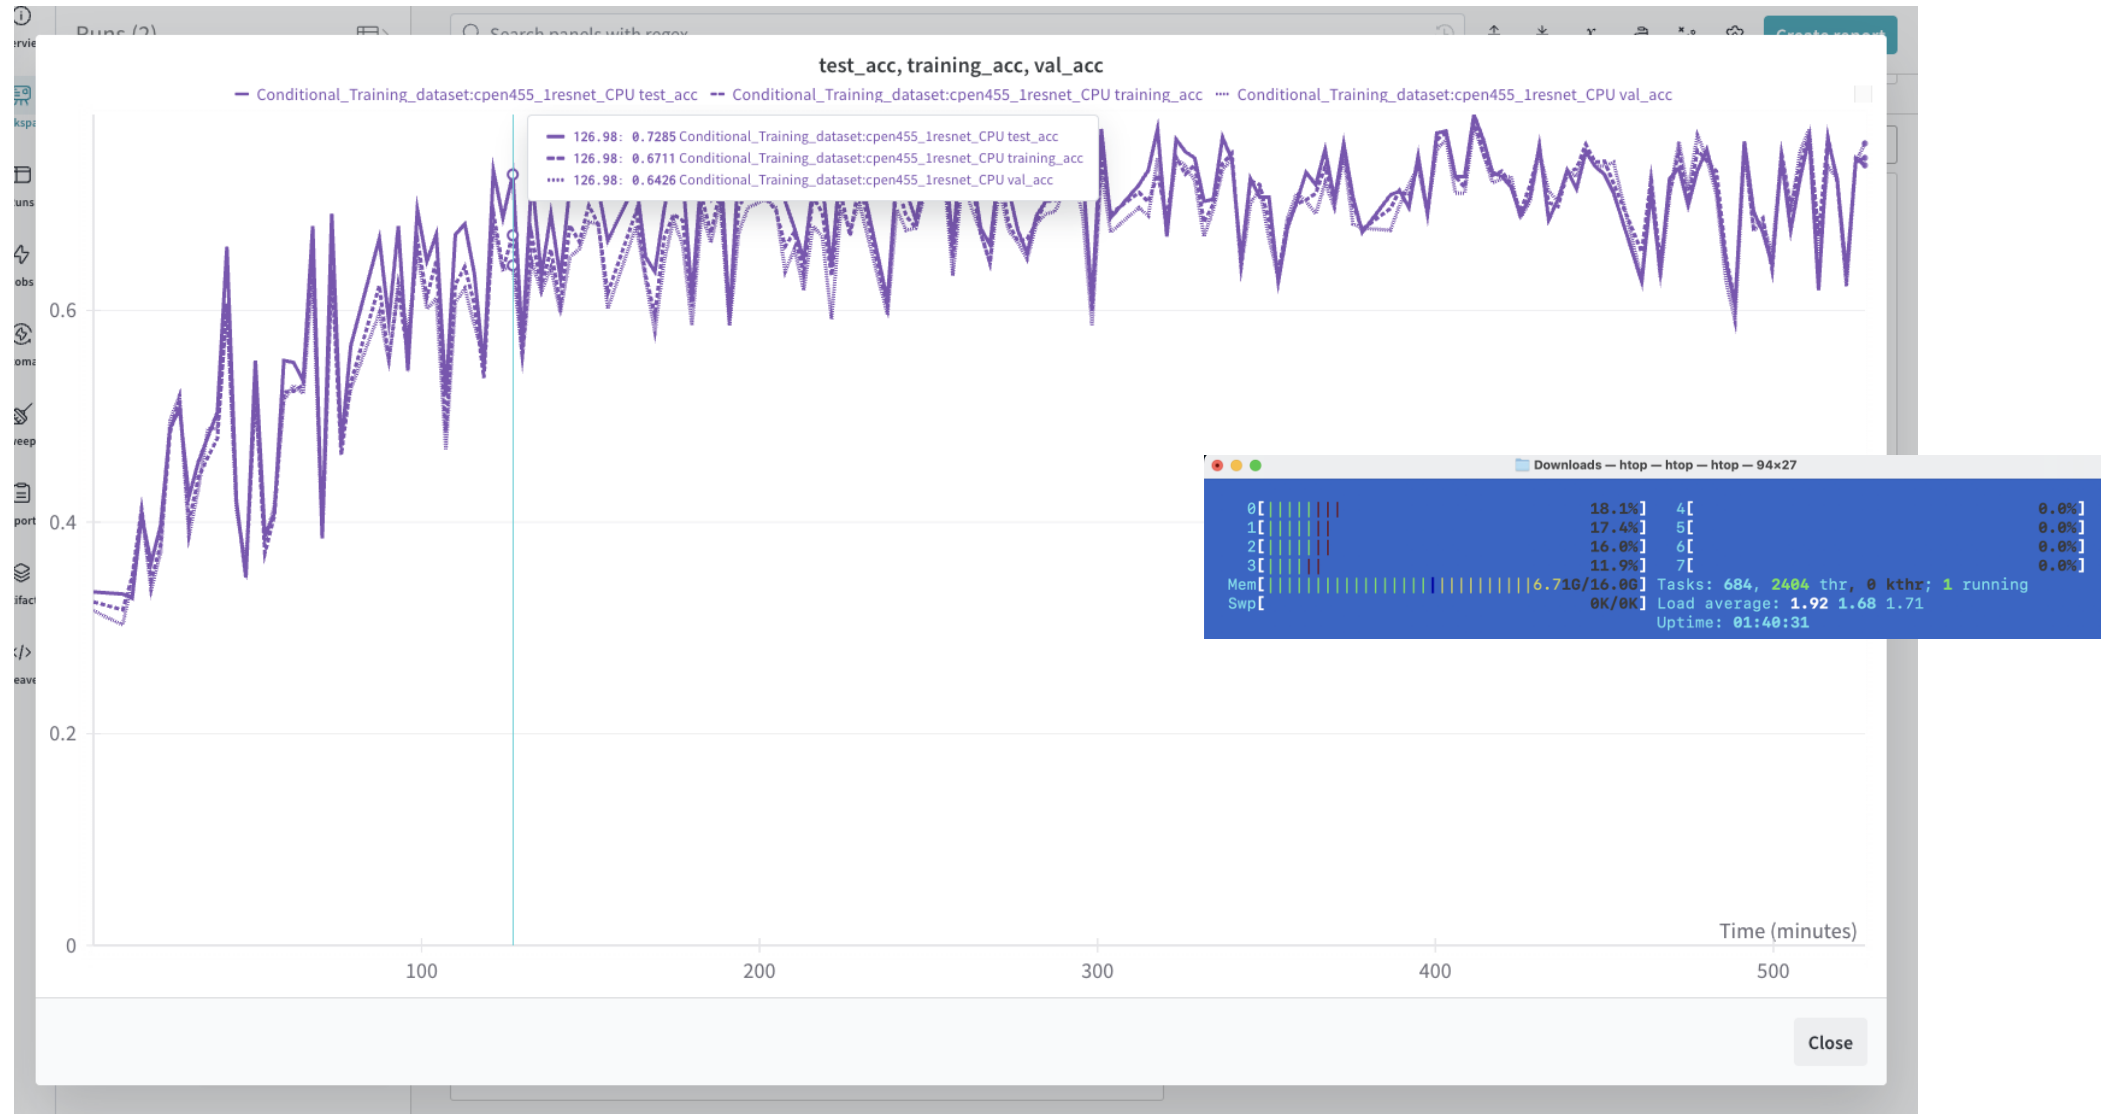
$$P(i|x) = \exp(\log P(i|x))$$

Important Hyperparameters

```
dataset.py  pcnn_train.py ×  utils.py  model.py
pcnn_train.py > ...
67
68  # model
69  parser.add_argument('-q', '--nr_resnet', type=int, default=5,
70                      help='Number of residual blocks per stage of the model')
71  parser.add_argument('-n', '--nr_filters', type=int, default=160,
72                      help='Number of filters to use across the model. Higher = larger model.')
73  parser.add_argument('-m', '--nr_logistic_mix', type=int, default=10,
74                      help='Number of logistic components in the mixture. Higher = more flexible model')
```


How long it takes to train a conditional PixelCNN++?

<https://wandb.ai/qihangz-work/CPEN455HW?nw=nwuserqihangz>



Interfaces

- `Generation_evaluation.py`
 - Fréchet Inception Distance
- `Classification_evaluation.py`
 - Validation set accuracy
 - Note that test set is for the final grading

Takeaways

- The logistic mixture model is **more stable** and achieves higher log-likelihood scores. It captures pixel dependencies and results in **smoother, more natural image generation**.
- Instead of using a fully autoregressive approach, PCNN++ implements two-stream processing, improving spatial coherence and long-range dependencies.
- Use downsampling to efficiently capture structure at multiple resolutions. Use a **coarse-to-fine hierarchical approach**, modeling a lower-resolution version first, then refining details.

Questions?