

F. TensorFlow codes

F.1. Probabilistic model and loss function

```
1 import tensorflow as tf
2
3 # inputs and annotators' labels
4 images = tf.placeholder(tf.float32, (None, image_shape))
5 labels = tf.placeholder(tf.int32, (None, num_annotators, num_classes))
6
7 # classifier for estimating true label distribution
8 logits = classifier(images)
9
10 # confusion matrices of annotators
11 confusion_matrices = confusion_matrix_estimators(num_annotators, num_classes)
12
13 # loss function
14 # 1. weighted cross-entropy
15 weighted_cross_entropy = cross_entropy_over_annotators(labels, logits, confusion_matrices)
16
17 # 2. trace of confusion matrices:
18 trace_norm = tf.reduce_mean(tf.trace(confusion_matrices))
19
20 # final loss (eq.(4))
21 total_loss = weighted_cross_entropy + scale * trace_norm
```

Listing 1: Implementation of the probabilistic model and the proposed loss function given in eq. (4). The final loss is minimized to learn jointly the confusion matrices of the respective annotators and the parameters of the classifier. The details of the used functions are given in Sec. F.2 & F.3

F.2. Defining confusion matrices of annotators

```
1 import numpy as np
2 import tensorflow as tf
3
4 def confusion_matrix_estimators(num_annotators, num_classes):
5     """Defines confusion matrix estimators.
6     This function defines a set of confusion matrices that characterize respective annotators.
7     Here (i, j)th element in the annotator confusion matrix of annotator a is given by
8     P(label_annotator_a = j | label_true = i) i.e. the probability that the annotator assigns label j to
9     the image when the ground truth label is i.
10
11     Args:
12         num_annotators: Number of annotators
13         num_classes: Number of classes.
14
15     Returns:
16         confusion_matrices: Annotator confusion matrices. A 'Tensor' of shape of shape [num_annotators,
17         num_classes, num_classes]
18     """
19     # initialise so the confusion matrices are close to identity matrices.
20     w_init = tf.constant(
21         np.stack([6.0 * np.eye(num_classes) - 5.0 for j in range(num_annotators)]),
22         dtype=tf.float32,
23     )
24     rho = tf.Variable(w_init, name='rho')
25
26     # ensure positivity
27     rho = tf.nn.softplus(rho)
28
29     # ensure each row sums to one
30     confusion_matrices = tf.divide(rho, tf.reduce_sum(rho, axis=-1, keepdims=True))
31     return confusion_matrices
```

Listing 2: The confusion matrices are defined as `tf.Variable`. The positivity of the elements of confusion matrices is ensured by passing them through a soft-plus function.

F.3. Cross-entropy loss with sparse and noisy labels

```
1 import numpy as np
2 import tensorflow as tf
3
4 def cross_entropy_over_annotators(labels, logits, confusion_matrices):
5     """ Cross entropy between noisy labels from multiple annotators and their confusion matrix models.
6     Args:
7         labels: One-hot representation of labels from multiple annotators.
8             tf.Tensor of size [batch, num_annotators, num_classes]. Missing labels are assumed to be
9             represented as zero vectors.
10        logits: Logits from the classifier. tf.Tensor of size [batch, num_classes]
11        confusion_matrices: Confusion matrices of annotators. tf.Tensor of size
12            [num_annotators, num_classes, num_classes]. The (i, j) th element of the confusion matrix
13            for annotator a denotes the probability  $P(\text{label\_annotator\_a} = j | \text{label\_true} = i)$ .
14    Returns:
15        The average cross-entropy across annotators and image examples.
16    """
17    # Treat one-hot labels as probability vectors
18    labels = tf.cast(labels, dtype=tf.float32)
19
20    # Sequentially compute the loss for each annotator
21    losses_all_annotators = []
22    for idx, labels_annotator in enumerate(tf.unstack(labels, axis=1)):
23        loss = sparse_confusion_matrix_softmax_cross_entropy(
24            labels=labels_annotator,
25            logits=logits,
26            confusion_matrix=confusion_matrices[idx, :, :],
27        )
28        losses_all_annotators.append(loss)
29
30    # Stack them into a tensor of size (batch, num_annotators)
31    losses_all_annotators = tf.stack(losses_all_annotators, axis=1)
32
33    # Filter out annotator networks with no labels. This allows you train
34    # annotator networks only when the labels are available.
35    has_labels = tf.reduce_sum(labels, axis=2)
36    losses_all_annotators = losses_all_annotators * has_labels
37    return tf.reduce_mean(tf.reduce_sum(losses_all_annotators, axis=1))
38
39 def sparse_confusion_matrix_softmax_cross_entropy(labels, logits, confusion_matrix):
40     """Cross entropy between noisy labels and confusion matrix based model for a single annotator.
41     Args:
42         labels: One-hot representation of labels. Tensor of size [batch, num_classes].
43         logits: Logits from the classifier. Tensor of size [batch, num_classes]
44         confusion_matrix: Confusion matrix of the annotator. Tensor of size [num_classes, num_classes].
45     Returns:
46         The average cross-entropy across annotators for image examples
47         Returns a `Tensor` of size [batch_size].
48     """
49    # get the predicted label distribution
50    preds_true = tf.nn.softmax(logits)
51
52    # Map label distribution into annotator label distribution by
53    # multiplying it by its confusion matrix.
54    preds_annotator = tf.matmul(preds_true, confusion_matrix)
55
56    # cross entropy
57    preds_clipped = tf.clip_by_value(preds_annotator, 1e-10, 0.9999999)
58    cross_entropy = tf.reduce_sum(-labels * tf.log(preds_clipped), axis=-1)
59    return cross_entropy
```

Listing 3: Implementation of the cross entropy loss function. Here we use zero vectors as the “one-hot” representations of missing labels.