
Supplementary Material: Toward Robustness against Label Noise in Training Deep Discriminative Neural Networks

Arash Vahdat
D-Wave Systems Inc.
Burnaby, BC, Canada
avahdat@dwavesys.com

1 Visualization

As shown in the E step in Sec. 3.3, the variational distribution q infers latent clean labels by combining information from both the image-based CRF-CNN model $p_{\theta}(\hat{\mathbf{y}}, \mathbf{h}|\mathbf{y}, \mathbf{x})$ and the label-based auxiliary distribution $p_{aux}^{\alpha}(\hat{\mathbf{y}}, \mathbf{h}|\mathbf{y})$. In our experiments, we observe that in general q proposes clean labels more accurately than the auxiliary distribution. Fig. 1 compares q against p_{aux} in terms of its ability to infer clean labels for a few instances in the noisy training set (D_N) for the COCO experiment with actual Flickr tags. In Fig. 2, examples of the recovered clean labels are visualized for the CIFAR-10 experiment.

	<p>Flickr \emptyset</p> <p>p_{aux} person</p> <p>q skateboard, person</p> <p>clean skateboard, person</p>
	<p>Flickr \emptyset</p> <p>p_{aux} person</p> <p>q person, baseball glove, baseball bat</p> <p>clean person, baseball glove, baseball bat, sports ball, chair, bench</p>
	<p>Flickr 2009, miami</p> <p>p_{aux} person</p> <p>q person, tennis racket</p> <p>clean person, tennis racket</p>
	<p>Flickr uploaded:by=flickr_mobile, flickriosapp:filter=NoFilter</p> <p>p_{aux} person</p> <p>q person, surfboard</p> <p>clean person, surfboard</p>
	<p>Flickr computer</p> <p>p_{aux} \emptyset</p> <p>q laptop, mouse, tv, keyboard</p> <p>clean laptop, mouse, tv, keyboard</p>
	<p>Flickr square, iphoneography, square format, instagram app, uploaded:by=instagram</p> <p>p_{aux} \emptyset</p> <p>q cup, dining table, bottle, bowl</p> <p>clean cup, dining table, bottle, bowl, spoon, hot dog</p>
	<p>Flickr food, square, square format, nikon, white, orange, fruit, color, India, photography, table, project365, bowl, colour, wood, 50mm, nikkor, bokeh</p> <p>p_{aux} orange, apple, banana</p> <p>q orange, bowl, dining table</p> <p>clean orange, bowl</p>
	<p>Flickr home, light, photo, art, chair, room, table, architecture, apartment, interior, couch, decor, beauty, design, live, lamp, indoor, furniture, relaxed, sofa, flooring, modern</p> <p>p_{aux} chair, couch, vase, book, dining table, sink, clock, bed, potted plant</p> <p>q chair, couch, vase, book</p> <p>clean chair, dining table, tv</p>

Figure 1: Visualization of inferred labels for a few instances in the noisy training set (D_N) of the COCO dataset. Flickr labels represent the noisy labels extracted from Flickr tags, whereas clean labels are the true labels ignored during training. p_{aux} and q correspond to the labels that are extracted using these distribution by thresholding at 0.5. The auxiliary distribution p_{aux} tends to assign the label “person” to the images with no tag while q adds more clean labels. In the last two images, q removes a few unrelated labels.



(a) cat \rightarrow dog



(b) dog \rightarrow cat



(c) automobile \rightarrow truck



(d) horse \rightarrow deer

Figure 2: Our proposed model can recover clean labels in the noisy training dataset. Here, corrupted instances are visualized for different categories in the CIFAR-10 training dataset. Sub-figures (a) through (d), captioned with *annotated label* \rightarrow *inferred label*, represents the instances that are labeled with the annotated label but have been assigned to the inferred label by our proposed variational distribution q . In this visualization, images are sorted based on the confidence of q for the inferred label from left to right and top to bottom, and the mistaken instances are marked with the red frame. The probability that q assigns for the inferred label is typically very high (> 0.9) for these images, which indicates that q is confident in changing the noisy labels.