

A Appendix

A.1 AutoAugment Controller Training Details

Table 6. Table of all the possible transformations that can be applied to an image. These are the transformations that are available to the controller during the search process. The range of magnitudes that the controller can predict for each of the transforms is listed in the third column. Some transformations do not have a magnitude associated with them (e.g. Equalize).

| Operation Name | Description | Range of magnitudes |
|----------------|---|---------------------|
| ShearX(Y) | Shear the image and the corners of the bounding boxes along the horizontal (vertical) axis with rate <i>magnitude</i> . | [-0.3,0.3] |
| TranslateX(Y) | Translate the image and the bounding boxes in the horizontal (vertical) direction by <i>magnitude</i> number of pixels. | [-150,150] |
| Rotate | Rotate the image and the bounding boxes <i>magnitude</i> degrees. | [-30,30] |
| Equalize | Equalize the image histogram. | |
| Solarize | Invert all pixels above a threshold value of <i>magnitude</i> . | [0,256] |
| SolarizeAdd | For each pixel in the image that is less than 128, add an additional amount to it decided by the magnitude. | [0,110] |
| Contrast | Control the contrast of the image. A <i>magnitude</i> =0 gives a gray image, whereas <i>magnitude</i> =1 gives the original image. | [0.1,1.9] |
| Color | Adjust the color balance of the image, in a manner similar to the controls on a colour TV set. A <i>magnitude</i> =0 gives a black & white image, whereas <i>magnitude</i> =1 gives the original image. | [0.1,1.9] |
| Brightness | Adjust the brightness of the image. A <i>magnitude</i> =0 gives a black image, whereas <i>magnitude</i> =1 gives the original image. | [0.1,1.9] |
| Sharpness | Adjust the sharpness of the image. A <i>magnitude</i> =0 gives a blurred image, whereas <i>magnitude</i> =1 gives the original image. | [0.1,1.9] |
| Cutout [6, 46] | Set a random square patch of side-length <i>magnitude</i> pixels to gray. | [0,60] |
| BBox_Only_X | Apply X to each bounding box content with independent probability, and magnitude that was chosen for X above. Location and the size of the bounding box are not changed. | |

Table 7. The sub-policies used in our learned augmentation policy. P and M correspond to the probability and magnitude with which the operations were applied in the sub-policy. Note that for each image in each mini-batch, one of the sub-policies is picked uniformly at random. The *No operation* is listed when an operation has a learned probability or magnitude of 0

| Operation 1 | P | M | Operation 2 | P | M |
|-----------------------------------|-----|----|----------------------|-----|----|
| Sub-policy 1 TranslateX | 0.6 | 4 | Equalize | 0.8 | 10 |
| Sub-policy 2 BBox_Only_TranslateY | 0.2 | 2 | Cutout | 0.8 | 8 |
| Sub-policy 3 ShearY | 1.0 | 2 | BBox_Only_TranslateY | 0.6 | 6 |
| Sub-policy 4 Rotate | 0.6 | 10 | Color | 1.0 | 6 |
| Sub-policy 5 No operation | | | No operation | | |