Module 8.5 : Dataset augmentation

Different forms of regularization

- l_2 regularization
- Dataset augmentation
- Parameter Sharing and tying
- Adding Noise to the inputs
- Adding Noise to the outputs
- Early stopping
- Ensemble methods
- Dropout

Different forms of regularization

- l_2 regularization
- Dataset augmentation
- Parameter Sharing and tying
- Adding Noise to the inputs
- Adding Noise to the outputs
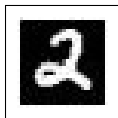- Early stopping
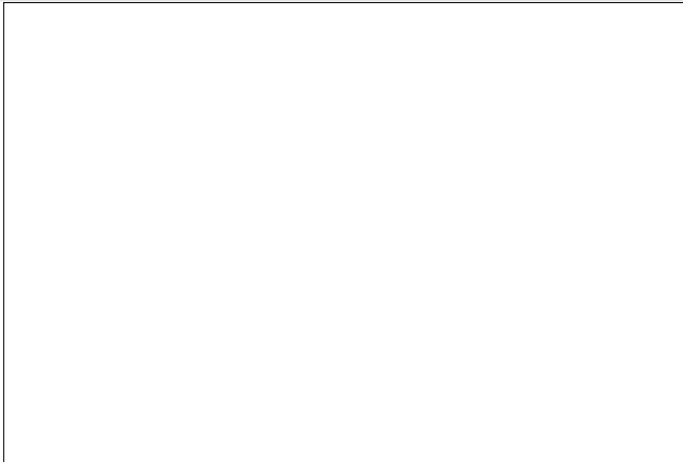- Ensemble methods
- Dropout

label = 2
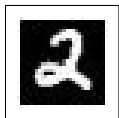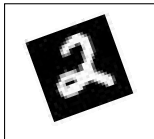
label = 2

[given training data]

label = 2

[given training data]

label = 2

[given training data]

rotated by 20°

rotated by 20°   rotated by 65°
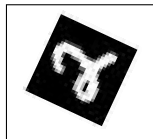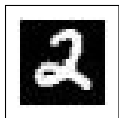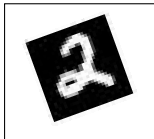

label = 2

[given training data]

rotated by 20°   rotated by 65°   shifted vertically

label = 2

[given training data]

label = 2

[given training data]

rotated by 20°   rotated by 65°   shifted vertically

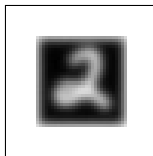shifted horizontally

label = 2

[given training data]

rotated by 20°　　rotated by 65°　　shifted vertically

shifted horizontally　　blurred

label = 2

[given training data]

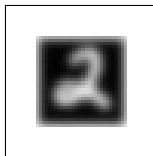rotated by 20°   rotated by 65°   shifted vertically

shifted horizontally   blurred   changed some pixels

label = 2

[given training data]

rotated by 20°   rotated by 65°   shifted vertically

shifted horizontally   blurred   changed some pixels
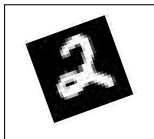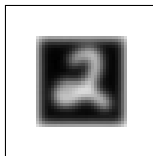
label = 2

label = 2

[given training data]

rotated by 20°   rotated by 65°   shifted vertically

shifted horizontally   blurred   changed some pixels

label = 2

[augmented data = created using some knowledge of the task]

label = 2

[given training data]
We exploit the fact that
certain transformations
to the image do not
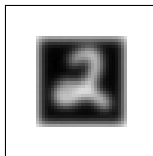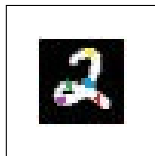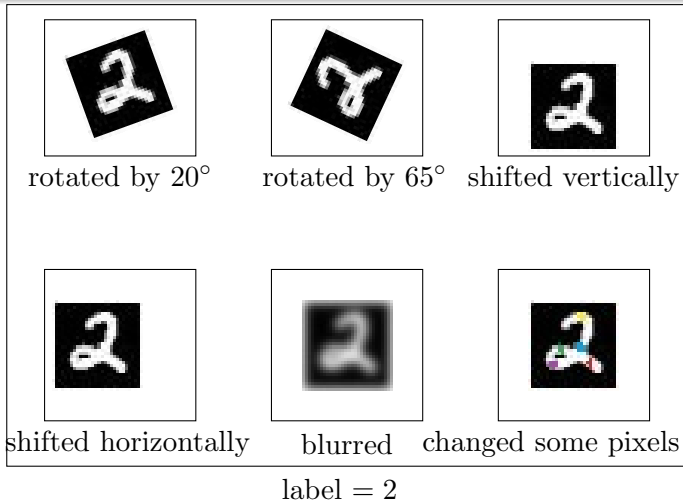change the label of the
image.

rotated by 20°   rotated by 65°   shifted vertically

shifted horizontally   blurred   changed some pixels

label = 2

[augmented data = created using some knowledge of the task]

- Typically, More data = better learning

- Typically, More data = better learning
- Works well for image classification / object recognition tasks

- Typically, More data = better learning
- Works well for image classification / object recognition tasks
- Also shown to work well for speech

- Typically, More data = better learning
- Works well for image classification / object recognition tasks
- Also shown to work well for speech
- For some tasks it may not be clear how to generate such data